

# DWARF Debugging Information Format Version 6



DWARF Debugging Information Format  
Committee

<http://www.dwarfstd.org>

February 27, 2024

\*\*\*WORKING DRAFT\*\*\*

# Copyright

DWARF Debugging Information Format, Version 6

Copyright © 2010, 2017, 2022, 2023 DWARF Debugging Information Format  
Committee

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.3; with no Invariant Sections, with no Front-Cover Texts, and with no Back-Cover Texts.

A copy of the license is included in the section entitled “GNU Free Documentation License.”

This document is based in part on the DWARF Debugging Information Format, Version 2, which contained the following notice:

UNIX International

Programming Languages SIG

Revision: 2.0.0 (July 27, 1993)

Copyright © 1992, 1993 UNIX International, Inc.

Permission to use, copy, modify, and distribute this documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appears in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name UNIX International not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission. UNIX International makes no representations about the suitability of this documentation for any purpose. It is provided “as is” without express or implied warranty.

This document is further based on the DWARF Debugging Information Format, Version 3 and Version 4, which are subject to the GNU Free Documentation License.

Trademarks:

- Intel386 is a trademark of Intel Corporation.
- Java is a trademark of Oracle, Inc.
- All other trademarks found herein are property of their respective owners.

# Foreword

The DWARF Debugging Information Format Committee was originally organized in 1988 as the Programming Languages Special Interest Group (PLSIG) of Unix International, Inc., a trade group organized to promote Unix System V Release 4 (SVR4).

PLSIG drafted a standard for DWARF Version 1, compatible with the DWARF debugging format used at the time by SVR4 compilers and debuggers from AT&T. This was published as Revision 1.1.0 on October 6, 1992. PLSIG also designed the DWARF Version 2 format, which followed the same general philosophy as Version 1, but with significant new functionality and a more compact, though incompatible, encoding. An industry review draft of DWARF Version 2 was published as Revision 2.0.0 on July 27, 1993.

Unix International dissolved shortly after the draft of Version 2 was released; no industry comments were received or addressed, and no final standard was released. The committee mailing list was hosted by OpenGroup (formerly XOpen).

The Committee reorganized in October, 1999, and met for the next several years to address issues that had been noted with DWARF Version 2 as well as to add a number of new features. In mid-2003, the Committee became a workgroup under the Free Standards Group (FSG), an industry consortium chartered to promote open standards. DWARF Version 3 was published on December 20, 2005, following industry review and comment.

The DWARF Committee withdrew from the Free Standards Group in February, 2007, when FSG merged with the Open Source Development Labs to form The Linux Foundation, more narrowly focused on promoting Linux. The DWARF Committee has been independent since that time.

It is the intention of the DWARF Committee that migrating from an earlier version of the DWARF standard to the current version should be straightforward and easily accomplished. Almost all constructs from DWARF Version 2 onward have been retained unchanged in DWARF Version 6, although a few have been compatibly superseded by improved constructs which are more compact and/or more expressive.

This document was created using the L<sup>A</sup>T<sub>E</sub>X document preparation system.

## The DWARF Debugging Information Format Committee

The DWARF Debugging Information Format Committee is open to compiler and debugger developers who have experience with source language debugging and debugging formats, and have an interest in promoting or extending the DWARF debugging format.

DWARF Committee members contributing to Version 6 are:

Todd Allen	Concurrent Real-Time
Pedro Alves	Pedro Alves Services
David Anderson, Associate Editor	
David Blaikie	Google
Ron Brender, Editor	
Andrew Cagney	
Eric Christopher	Google
Cary Coutant, Chair (from March 2023)	
John DelSignore	Perforce
Jonas Devlieghere	Apple
Michael Eager, past Chair (to February 2023)	Eager Consulting
Jini Susan George	AMD
Tommy Hoffner	Untether AI
Jakub Jelínek	Red Hat
Simon Marchi	EfficiOS
Jason Merrill	Red Hat
Markus Metzger	Intel
Jeremy Morse	Sony
Adrian Prantl	Apple
Hafiz Abid Qadeer	Mentor Graphics
Paul Robinson	Sony
Tom Russell	Sony
Fāng-ruì Sòng	Google
Caroline Tice	Google
Tom Tromeey	Adacore
Tony Tye	AMD
Keith Walker	Arm
Mark Wielaard	Red Hat
Brock Wyma	Intel
Jian Xu	IBM
Zoran Zaric	AMD

For further information about DWARF or the DWARF Committee, see:

<http://www.dwarfstd.org>

## How to Use This Document

This document is intended to be usable in online as well as traditional paper forms. Both online and paper forms include page numbers, a Table of Contents, a List of Figures, a List of Tables and an Index.

Text in normal font describes required aspects of the DWARF format. Text in *italics* is explanatory or supplementary material, and not part of the format definition itself.

### *Online Form*

In the online form, [blue text](#) is used to indicate hyperlinks. Most hyperlinks link to the definition of a term or construct, or to a cited Section or Figure. However, attributes in particular are often used in more than one way or context so that there is no single definition; for attributes, hyperlinks link to the introductory table of all attributes which in turn contains hyperlinks for the multiple usages.

The occurrence of a DWARF name in its definition (or one of its definitions in the case of some attributes) is shown in **red text**. Other occurrences of the same name in the same or possibly following paragraphs are generally in normal text color.)

The Table of Contents, List of Figures, List of Tables and Index provide hyperlinks to the respective items and places.

### *Paper Form*

In the traditional paper form, the appearance of the hyperlinks and definitions on a page of paper does not distract the eye because the blue hyperlinks and the color used for definitions are typically imaged by black and white printers in a manner nearly indistinguishable from other text. (Hyperlinks are not underlined for this same reason.)



# Contents

<b>Contents</b>	<b>vii</b>
<b>List of Figures</b>	<b>xii</b>
<b>List of Tables</b>	<b>xvi</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Purpose and Scope . . . . .	1
1.2 Overview . . . . .	2
1.3 Objectives and Rationale . . . . .	2
1.4 Changes from Version 5 to Version 6 . . . . .	8
1.5 Changes from Version 4 to Version 5 . . . . .	8
1.6 Changes from Version 3 to Version 4 . . . . .	10
1.7 Changes from Version 2 to Version 3 . . . . .	11
1.8 Changes from Version 1 to Version 2 . . . . .	12
<b>2 General Description</b>	<b>15</b>
2.1 The Debugging Information Entry (DIE) . . . . .	15
2.2 Attribute Types . . . . .	17
2.3 Relationship of Debugging Information Entries . . . . .	25
2.4 Target Addresses . . . . .	26
2.5 DWARF Expressions . . . . .	26
2.6 Location Descriptions . . . . .	38
2.7 Types of Program Entities . . . . .	47
2.8 Accessibility of Declarations . . . . .	47
2.9 Visibility of Declarations . . . . .	47
2.10 Virtuality of Declarations . . . . .	48
2.11 Artificial Entries . . . . .	48
2.12 Address Classes . . . . .	49
2.13 Non-Defining Declarations and Completions . . . . .	49
2.14 Declaration Coordinates . . . . .	50

# CONTENTS

2.15	Identifier Names . . . . .	51
2.16	Data Locations and DWARF Procedures . . . . .	51
2.17	Code Addresses, Ranges and Base Addresses . . . . .	52
2.18	Entry Address . . . . .	56
2.19	Static and Dynamic Values of Attributes . . . . .	56
2.20	Entity Descriptions . . . . .	57
2.21	Byte and Bit Sizes . . . . .	58
2.22	Linkage Names . . . . .	58
2.23	Template Parameters . . . . .	58
2.24	Alignment . . . . .	59
<b>3</b>	<b>Program Scope Entries</b>	<b>61</b>
3.1	Unit Entries . . . . .	61
3.2	Module, Namespace and Importing Entries . . . . .	73
3.3	Subroutine and Entry Point Entries . . . . .	78
3.4	Call Site Entries and Parameters . . . . .	91
3.5	Lexical Block Entries . . . . .	95
3.6	Label Entries . . . . .	96
3.7	With Statement Entries . . . . .	96
3.8	Try and Catch Block Entries . . . . .	97
3.9	Declarations with Reduced Scope . . . . .	97
<b>4</b>	<b>Data Object and Object List</b>	<b>99</b>
4.1	Data Object Entries . . . . .	99
4.2	Common Block Entries . . . . .	102
4.3	Namelist Entries . . . . .	103
<b>5</b>	<b>Type Entries</b>	<b>106</b>
5.1	Base Type Entries . . . . .	106
5.2	Unspecified Type Entries . . . . .	112
5.3	Type Modifier Entries . . . . .	113
5.4	Typedef Entries . . . . .	115
5.5	Array Type Entries . . . . .	115
5.6	Coarray Type Entries . . . . .	117
5.7	Structure, Union, Class and Interface Type Entries . . . . .	118
5.8	Condition Entries . . . . .	129
5.9	Enumeration Type Entries . . . . .	129
5.10	Subroutine Type Entries . . . . .	131
5.11	String Type Entries . . . . .	132
5.12	Set Type Entries . . . . .	133
5.13	Subrange Type Entries . . . . .	133



# CONTENTS

5.14	Pointer to Member Type Entries . . . . .	135
5.15	File Type Entries . . . . .	136
5.16	Dynamic Type Entries . . . . .	136
5.17	Template Alias Entries . . . . .	137
5.18	Dynamic Properties of Types . . . . .	137
<b>6</b>	<b>Other Debugging Information</b>	<b>141</b>
6.1	Accelerated Access . . . . .	141
6.2	Line Number Information . . . . .	155
6.3	Macro Information . . . . .	172
6.4	Call Frame Information . . . . .	178
<b>7</b>	<b>Data Representation</b>	<b>191</b>
7.1	Vendor Extensibility . . . . .	191
7.2	Reserved Values . . . . .	192
7.3	Relocatable, Split, Executable, Shared, Package and Supplementary Object Files . . . . .	193
7.4	32-Bit and 64-Bit DWARF Formats . . . . .	204
7.5	Format of Debugging Information . . . . .	208
7.6	Variable Length Data . . . . .	232
7.7	DWARF Expressions and Location Descriptions . . . . .	233
7.8	Base Type Attribute Encodings . . . . .	238
7.9	Accessibility Codes . . . . .	240
7.10	Visibility Codes . . . . .	240
7.11	Virtuality Codes . . . . .	241
7.12	Source Languages . . . . .	241
7.13	Address Class Encodings . . . . .	243
7.14	Identifier Case . . . . .	243
7.15	Calling Convention Encodings . . . . .	244
7.16	Inline Codes . . . . .	244
7.17	Array Ordering . . . . .	244
7.18	Discriminant Lists . . . . .	245
7.19	Name Index Table . . . . .	245
7.20	Defaulted Member Encodings . . . . .	246
7.21	Address Range Table . . . . .	246
7.22	Line Number Information . . . . .	247
7.23	Macro Information . . . . .	249
7.24	Call Frame Information . . . . .	251
7.25	Range List Entries for Non-contiguous Address Ranges . . . . .	252
7.26	String Offsets Table . . . . .	253
7.27	Address Table . . . . .	253

# CONTENTS

7.28	Range List Table . . . . .	254
7.29	Location List Table . . . . .	255
7.30	Dependencies and Constraints . . . . .	256
7.31	Integer Representation Names . . . . .	257
7.32	Type Signature Computation . . . . .	257
7.33	Name Table Hash Function . . . . .	262
7.34	Contiguous Tables . . . . .	262
<b>A</b>	<b>Attributes by Tag (Informative)</b>	<b>263</b>
<b>B</b>	<b>Debug Section Relationships (Informative)</b>	<b>284</b>
B.1	Normal DWARF Section Relationships . . . . .	284
B.2	Split DWARF Section Relationships . . . . .	285
<b>C</b>	<b>Encoding/Decoding (Informative)</b>	<b>294</b>
<b>D</b>	<b>Examples (Informative)</b>	<b>298</b>
D.1	General Description Examples . . . . .	298
D.2	Aggregate Examples . . . . .	304
D.3	Namespace Examples . . . . .	331
D.4	Member Function Examples . . . . .	335
D.5	Line Number Examples . . . . .	339
D.6	Call Frame Information Example . . . . .	343
D.7	Inlining Examples . . . . .	347
D.8	Constant Expression Example . . . . .	356
D.9	Unicode Character Example . . . . .	358
D.10	Type-Safe Enumeration Example . . . . .	359
D.11	Template Examples . . . . .	360
D.12	Template Alias Examples . . . . .	362
D.13	Implicit Pointer Examples . . . . .	365
D.14	String Type Examples . . . . .	369
D.15	Call Site Examples . . . . .	371
D.16	Macro Example . . . . .	379
D.17	Parameter Default Value Examples . . . . .	383
<b>E</b>	<b>Compression (Informative)</b>	<b>386</b>
E.1	Using Compilation Units . . . . .	386
E.2	Using Type Units . . . . .	396
E.3	Summary of Compression Techniques . . . . .	409
<b>F</b>	<b>Split DWARF Object Files (Informative)</b>	<b>412</b>
F.1	Overview . . . . .	412

## CONTENTS

F.2	Split DWARF Object File Example . . . . .	417
F.3	DWARF Package File Example . . . . .	430
<b>G</b>	<b>Section Version Numbers (Informative)</b>	<b>436</b>
<b>H</b>	<b>GNU Free Documentation License</b>	<b>440</b>
H.1	APPLICABILITY AND DEFINITIONS . . . . .	441
H.2	VERBATIM COPYING . . . . .	442
H.3	COPYING IN QUANTITY . . . . .	443
H.4	MODIFICATIONS . . . . .	444
H.5	COMBINING DOCUMENTS . . . . .	446
H.6	COLLECTIONS OF DOCUMENTS . . . . .	446
H.7	AGGREGATION WITH INDEPENDENT WORKS . . . . .	446
H.8	TRANSLATION . . . . .	447
H.9	TERMINATION . . . . .	447
H.10	FUTURE REVISIONS OF THIS LICENSE . . . . .	448
H.11	RELICENSING . . . . .	448
	<b>Index</b>	<b>452</b>

# List of Figures

5.1	Type modifier examples . . . . .	114
6.1	Name Index Layout . . . . .	145
7.1	Name Table Hash Function Definition . . . . .	262
B.1	Debug section relationships . . . . .	286
B.2	Split DWARF section relationships . . . . .	290
C.1	Algorithm to encode an unsigned integer . . . . .	294
C.2	Algorithm to encode a signed integer . . . . .	295
C.3	Algorithm to decode an unsigned LEB128 integer . . . . .	295
C.4	Algorithm to decode a signed LEB128 integer . . . . .	296
D.1	Compilation units and abbreviations table . . . . .	299
D.2	Fortran array example: source fragment . . . . .	304
D.3	Fortran array example: descriptor representation . . . . .	305
D.4	Fortran array example: DWARF description . . . . .	308
D.5	Fortran scalar coarray: source fragment . . . . .	311
D.6	Fortran scalar coarray: DWARF description . . . . .	311
D.7	Fortran array coarray: source fragment . . . . .	311
D.8	Fortran array coarray: DWARF description . . . . .	311
D.9	Fortran multidimensional coarray: source fragment . . . . .	312
D.10	Fortran multidimensional coarray: DWARF description . . . . .	312
D.11	Declaration of a Fortran 2008 assumed-rank array . . . . .	313
D.12	One of many possible layouts for an array descriptor . . . . .	313
D.13	Sample DWARF for the array descriptor in Figure D.12 . . . . .	314
D.14	How to interpret the DWARF from Figure D.13 . . . . .	315
D.15	Fortran dynamic type example: source . . . . .	316
D.16	Fortran dynamic type example: DWARF description . . . . .	317
D.17	Anonymous structure example: source fragment . . . . .	318
D.18	Anonymous structure example: DWARF description . . . . .	318
D.19	Ada example: source fragment . . . . .	319

## LIST OF FIGURES

D.20 Ada example: DWARF description . . . . .	320
D.21 Packed record example: source fragment . . . . .	321
D.22 Packed record example: DWARF description . . . . .	321
D.23 Big-endian data bit offsets . . . . .	324
D.24 Little-endian data bit offsets . . . . .	324
D.25 Ada biased bit-field example: Ada source . . . . .	325
D.26 Ada biased bit-field example: DWARF description . . . . .	325
D.27 Pascal variant record example: source . . . . .	326
D.28 Pascal variant record example: DWARF description . . . . .	327
D.29 Ada variant record example: source . . . . .	328
D.30 Ada variant record example: DWARF description . . . . .	329
D.31 Rust enum example: source . . . . .	329
D.32 Rust enum example: DWARF description . . . . .	330
D.33 Namespace example #1: source fragment . . . . .	331
D.34 Namespace example #1: DWARF description . . . . .	332
D.35 Namespace example #2: source fragment . . . . .	334
D.36 Namespace example #2: DWARF description . . . . .	334
D.37 Member function example: source fragment . . . . .	335
D.38 Member function example: DWARF description . . . . .	335
D.39 Reference- and rvalue-reference-qualification example: source fragment	337
D.40 Reference- and rvalue-reference-qualification example: DWARF description . . . . .	338
D.41 Example line number program header . . . . .	339
D.42 Example line number special opcode mapping . . . . .	340
D.43 Line number program example: machine code . . . . .	341
D.44 Call frame information example: machine code fragments . . . . .	344
D.45 Inlining examples: pseudo-source fragment . . . . .	347
D.46 Inlining example #1: abstract instance . . . . .	349
D.47 Inlining example #1: concrete instance . . . . .	350
D.48 Inlining example #2: abstract instance . . . . .	352
D.49 Inlining example #2: concrete instance . . . . .	354
D.50 Inlining example #3: abstract instance . . . . .	355
D.51 Inlining example #3: concrete instance . . . . .	356
D.52 Constant expressions: C++ source . . . . .	356
D.53 Constant expressions: DWARF description . . . . .	357
D.54 Unicode character example: source . . . . .	358
D.55 Unicode character example: DWARF description . . . . .	358
D.56 Type-safe enumeration example: source . . . . .	359
D.57 Type-safe enumeration example: DWARF description . . . . .	359
D.58 C++ template example #1: source . . . . .	360
D.59 C++ template example #1: DWARF description . . . . .	360

## LIST OF FIGURES

D.60 C++ template example #2: source . . . . .	361
D.61 C++ template example #2: DWARF description . . . . .	361
D.62 C++ template alias example #1: source . . . . .	362
D.63 C++ template alias example #1: DWARF description . . . . .	363
D.64 C++ template alias example #2: source . . . . .	363
D.65 C++ template alias example #2: DWARF description . . . . .	364
D.66 C implicit pointer example #1: source . . . . .	365
D.67 C implicit pointer example #1: DWARF description . . . . .	366
D.68 C implicit pointer example #2: source . . . . .	367
D.69 C implicit pointer example #2: DWARF description . . . . .	368
D.70 String type example: source . . . . .	369
D.71 String type example: DWARF representation . . . . .	370
D.72 Call Site Example #1: Source . . . . .	371
D.73 Call Site Example #1: Code . . . . .	372
D.74 Call site example #1: DWARF encoding . . . . .	374
D.75 Call site example #2: source . . . . .	376
D.76 Call site example #2: code . . . . .	377
D.77 Call site example #2: DWARF encoding . . . . .	378
D.78 Macro example: source . . . . .	379
D.79 Macro example: simple DWARF encoding . . . . .	380
D.80 Macro example: sharable DWARF encoding . . . . .	381
D.81 Default value example #1: C++ source . . . . .	383
D.82 Default value example #1: DWARF encoding . . . . .	383
D.83 Default value example #2: Ada source . . . . .	384
D.84 Default value example #2: DWARF encoding . . . . .	384
E.1 Duplicate elimination example #1: C++ Source . . . . .	393
E.2 Duplicate elimination example #1: DWARF section group . . . . .	393
E.3 Duplicate elimination example #1: primary compilation unit . . . . .	394
E.4 Duplicate elimination example #2: Fortran source . . . . .	394
E.5 Duplicate elimination example #2: DWARF section group . . . . .	395
E.6 Duplicate elimination example #2: primary unit . . . . .	396
E.7 Duplicate elimination example #2: companion source . . . . .	396
E.8 Duplicate elimination example #2: companion DWARF . . . . .	397
E.9 Type signature examples: C++ source . . . . .	398
E.10 Type signature computation #1: DWARF representation . . . . .	399
E.11 Type signature computation #1: flattened byte stream . . . . .	400
E.12 Type signature computation #2: DWARF representation . . . . .	401
E.13 Type signature example #2: flattened byte stream . . . . .	403
E.14 Type signature example usage . . . . .	406
E.15 Type signature computation grammar . . . . .	407

## LIST OF FIGURES

E.16	Completing declaration of a member function: DWARF encoding . . .	408
F.1	Split object example: source fragment #1 . . . . .	417
F.2	Split object example: source fragment #2 . . . . .	418
F.3	Split object example: source fragment #3 . . . . .	419
F.4	Split object example: skeleton DWARF description . . . . .	420
F.5	Split object example: executable file DWARF excerpts . . . . .	422
F.6	Split object example: demo1.dwo excerpts . . . . .	424
F.7	Split object example: demo2.dwo DWARF .debug_info.dwo excerpts . .	427
F.8	Split object example: demo2.dwo DWARF .debug_loclists.dwo excerpts	429
F.9	Sections and contributions in example package file demo.dwp . . . . .	431
F.10	Example CU index section . . . . .	433
F.11	Example TU index section . . . . .	434

# List of Tables

2.1	Tag names	16
2.2	Attribute names	17
2.3	Classes of attribute value	23
2.4	Accessibility codes	47
2.5	Visibility codes	48
2.6	Virtuality codes	48
3.1	Language names	64
3.2	Version Encoding Schemes	66
3.3	Identifier case codes	67
3.4	Calling convention codes for subroutines	79
3.5	Inline codes	85
4.1	Endianness attribute values	102
5.1	Encoding attribute values	108
5.2	Decimal sign attribute values	111
5.3	Type modifier tags	113
5.4	Array ordering	116
5.5	Calling convention codes for types	120
5.6	Defaulted attribute names	126
5.7	Discriminant descriptor values	128
6.1	Index attribute encodings	153
6.3	State machine registers	157
6.4	Line number program initial state	159
7.1	DWARF package file section identifier encodings	202
7.2	Unit header unit type encodings	209
7.3	Tag encodings	214
7.4	Child determination encodings	217
7.5	Attribute encodings	218
7.6	Attribute form encodings	230



## LIST OF TABLES

7.7	Examples of unsigned LEB128 encodings	233
7.8	Examples of signed LEB128 encodings	233
7.9	DWARF operation encodings	234
7.10	Location list entry encoding values	238
7.11	Base type encoding values	238
7.12	Decimal sign encodings	240
7.13	Endianness encodings	240
7.14	Accessibility encodings	240
7.15	Visibility encodings	241
7.16	Virtuality encodings	241
7.17	Language encodings	242
7.18	Identifier case encodings	243
7.19	Calling convention encodings	244
7.20	Inline encodings	244
7.21	Ordering encodings	245
7.22	Discriminant descriptor encodings	245
7.23	Name index attribute encodings	245
7.24	Defaulted attribute encodings	246
7.25	Line number standard opcode encodings	247
7.26	Line number extended opcode encodings	248
7.27	Line number header entry format encodings	248
7.28	Macro information entry type encodings	250
7.29	Call frame instruction encodings	251
7.30	Range list entry encoding values	252
7.31	Integer representation names	257
7.32	Attributes used in type signature computation	259
A.1	Attributes by tag value	264
D.2	Line number program example: one encoding	342
D.3	Line number program example: alternate encoding	342
D.4	Call frame information example: conceptual matrix	344
D.5	Call frame information example: common information entry encoding	345
D.6	Call frame information example: frame description entry encoding	346
F.1	Unit attributes by unit kind	416
G.1	Section version numbers	437

## LIST OF TABLES

*(empty page)*

## Change Summary

# Change Summary

### Note

*This change summary is included only in draft versions of this document.*

<i>Date</i>	<i>Issue Incorporated or Other Change</i>
2/17/2021	<i>Begin DWARF Version 6. Update front matter.</i>
3/10/2021	<i>Remove change bars commands that were lingering from V5 (disabled in public release). Remove "New in DWARF Version 5" annotations.</i>
3/11/2021	<i>Issue 180613.1, stop using horizontal space to suppress ligatures.</i>
3/14/2021	<i>Issues 171130.1, 200505.1, 200505.2 and 200505.3, minor editorial corrections.</i>
3/23/2021	<i>Issues 200505.4 and 200505.7, editorial corrections. Issue 161206.2, add non-normative clarification re DW_OP_piece vs DW_OP_bit_piece.</i>
4/14/2021	<i>Remove 2005 from Copyright statement (was then the Free Standards Group).</i>
4/25/2021	<i>Issue 170527.1 re DW_IDX_external for external symbols.</i>
5/2/2021	<i>Start V6 column in version numbers appendix.</i>
5/3/2021	<i>Cleanup some table formatting in the L<sup>A</sup>T<sub>E</sub>X source.</i>
5/17/2021	<i>Issue 191025.1, DW_OP_bit_piece.</i>
5/21/2021	<i>Issue 180503.1, usage suggestions for LEB128 padding. Issue 170427.2, extending loclists.</i>
6/17/2021	<i>Issue 200427.1, missing link and related notes for Figure B.1, and Issue 200519.1, missing notes for Figure B.2. Issue 180426.2, add line number extended op DW_LNE_padding.</i>
6/30/2021	<i>180326.1, clarify consistency of DWARF 32/64 format within a CU.</i>
7/12/2021	<i>210218.1, index entry shows up in PDF.</i>
8/14/2021	<i>210628.1, clarification of relative paths in DW_AT_comp_dir. 200710.1, inconsistent description of data representation for the range list table.</i>
9/28/2021	<i>180625.1, inconsistent initial length descriptions. 181019.1, inconsistency in DW_AT_import descriptions.</i>
10/9/2021	<i>171103.1, DW_AT_call_origin should be encoded as reference class. 180426.1, Add DW_FORM_strp_sup to forms allowed in .debug_line vendor-defined content descriptions.</i>
10/30/2021	<i>200505.4, Augmentation string is null-terminated. See 3/23/2021. 200505.7, Declarations with reduced scope. See 3/23/2021 and 5/7/2022.</i>
11/21/2021	<i>200709.1, DW_AT_rnglists_base in DW_TAG_skeleton_unit 181205.1, Clarify DW_OP_piece documentation for parts of values that are optimized out.</i>

## Change Summary

<i>Date</i>	<i>Issue Incorporated or Other Change</i>
1/14/2022	200602.1, <i>.debug_macro.dwo</i> refers to <i>.debug_line.dwo</i> ? Also, tweak some member names and affiliations in the Foreword.
1/20/2022	210314.1, Eliminate all indefinite antecedents.
3/12/2022	210113.1, Allow zero-length entries in <i>.debug_aranges</i> . 200609.1, Reserve an address for "not present".
3/26/2022	201007.1, Wide registers in location description expressions. 210310.1, Clarify <i>DW_AT_rnglists_base</i> and <i>DW_FORM_rnglistx</i> in split DWARF. 210429.1, Clarify description of line number extended opcodes.
4/16/2022	180517.1, Variant parts without a discriminant. 210622.1, Typo in <i>.debug_rnglists</i> section header description.
5/7/2022	210208.2, Standardize <i>DW_AT_GNU_numerator</i> and <i>DW_AT_GNU_denominator</i> . 200505.4, Augmentation string. Reverses 10/30/2021.
5/30/2022	211101.1, Allow 64-bit string offsets in DWARF-32.
6/15/2022	210419.1, Split <i>DW_AT_language</i> into <i>DW_AT_language_name</i> and <i>DW_AT_language_version</i> .
7/5/2022	190809.1, Add <i>DW_AT_bias</i> .
7/17/2022	180201.1, Source text embedding.
8/6/2022	210713.1, Fix "file 0".
8/7/2022	211108.2, Allow non-uniform record formats in the file name table.
8/8/2022	211022.1, Empty range list entry. 181003.1, Forbid <i>DW_OP_call_ref</i> and <i>DW_FORM_addr_ref</i> in a <i>.dwo</i> file.
8/14/2022	220427.1, Deprecate the <i>DW_AT_segment</i> attribute.
9/4/2022	181223.1, Add Microsoft SourceLink support. 211108.2, Rework example in D.5 to illustrate <i>DW_LNCT_source</i> and <i>DW_LNCT_URL</i> . Review and adjust pagination.
10/12/2022	211108.2, Further rework of the example in D.5.
10/22/2022	211102.1, No <i>DW_FORM_strp</i> in <i>.dwo</i> files. 141117.1, Arbitrary expressions as formal parameter default values.
11/7/2022	220212.1, Disambiguate "ending address offset in location and range lists.
11/8/2022	211004.1, Replace <i>DW_MACRO_define/undefine_sup</i> with sized versions.
11/14/2022	220708.1, Remove edge (fo) from Figure B.2. 220711.1, Name Table index attribute. 220711.2, Name Table Figure 6.1.
11/14/2022	211103.1, Call site entries for optimized out functions.
11/30/2022	Incorporate minor review tweaks.
12/10/2022 <i>et al</i>	Additional minor review tweaks.
1/29/2023	210218.2, Generalize complex number support. 220708.2, <i>.debug_c,tu_index</i> missing/incomplete DWARF64 support.

## Change Summary

<i>Date</i>	<i>Issue Incorporated or Other Change</i>
4/3/2023	221031.1, Future-proof text from 211102.1. 220802.1, Introduce DW_FORM_addr_offset paired form. 170427.3, Extending loclists with common sublists. 220713.1, Name Table Figure 6.1. Update committee members list and roles.
6/15/2023	211108.1, Add DW_AT_artificial for DW_TAG_typedef. 220824.1, Use uniform encoding of DWARF expressions in CFI instructions.
6/27/2023	180123.1, Layout of discriminant entries in variant parts. 181026.3, Don't forbid extensions to the dwp file. 221118.1, Name Table 6.1.1.4.8.
7/10/2023	221114.1, DW_FORM_implicit_const and DW_FORM_indirect. 230223.1, Tidy up location description descriptions. 230414.1, Eliminate last use of "location expression".
8/6/2023	221203.1, Remove suggestion that DW_FORM_sec_offset may not be used for lists in split units.
10/24/2023	230103.1, Clarify that DW_CFA_remember_state includes the current CFA. 230120.1, DW_OP_call_ref & DW_OP_implicit_pointer correction. 230616.1, New form classes for values vs. location descriptions. 210514.1, Add GPU shading and kernel languages. 210115.1, DW_lang_code for the Netwide Assembler (NASM). 230203.1, C# language ID. 230502.1, New language name Mojo.
11/14/2023	230808.1, DW_OP_entry_value description. 230413.1, Tensor types.
12/3/2023	230329.1, Tables which have a unit_length header field must be contiguous. 230529.1, Bit-precise integer types.
1/15/2024	231230.1, New language code for Ruby. 231013.1, Tombstoning TU entries in .debug_names.
2/18/2024	230324.1, Expression operation vendor extensibility opcode. 230412.1, Ambiguity in Static and Dynamic Values of Attributes.

## LIST OF TABLES

# Chapter 1

## Introduction

This document defines a format for describing programs to facilitate user source level debugging. This description can be generated by compilers, assemblers and linkage editors. It can be used by debuggers and other tools. The debugging information format does not favor the design of any compiler or debugger. Instead, the goal is to create a method of communicating an accurate picture of the source program to any debugger in a form that is extensible to different languages while retaining compatibility.

The design of the debugging information format is open-ended, allowing for the addition of new debugging information to accommodate new languages or debugger capabilities while remaining compatible with other languages or different debuggers.

### 1.1 Purpose and Scope

The debugging information format described in this document is designed to meet the symbolic, source-level debugging needs of different languages in a unified fashion by requiring language independent debugging information whenever possible. Aspects of individual languages, such as C++ virtual functions or Fortran common blocks, are accommodated by creating attributes that are used only for those languages. This document is believed to cover most debugging information needs of Ada, C, C++, COBOL, and Fortran; it also covers the basic needs of various other languages.

This document describes DWARF Version 5, the fifth generation of debugging information based on the DWARF format. DWARF Version 5 extends DWARF Version 4 in a compatible manner.

## Chapter 1. Introduction

1 The intended audience for this document is the developers of both producers  
2 and consumers of debugging information, typically compilers, debuggers and  
3 other tools that need to interpret a binary program in terms of its original source.

### 4 **1.2 Overview**

5 There are two major pieces to the description of the DWARF format in this  
6 document. The first piece is the informational content of the debugging entries.  
7 The second piece is the way the debugging information is encoded and  
8 represented in an object file.

9 The informational content is described in Chapters 2 through 6. Chapter 2  
10 describes the overall structure of the information and attributes that are common  
11 to many or all of the different debugging information entries. Chapters 3, 4 and 5  
12 describe the specific debugging information entries and how they communicate  
13 the necessary information about the source program to a debugger. Chapter 6  
14 describes debugging information contained outside of the debugging  
15 information entries. The encoding of the DWARF information is presented in  
16 Chapter 7.

17 This organization closely follows that used in the DWARF Version 4 document.  
18 Except where needed to incorporate new material or to correct errors, the  
19 DWARF Version 4 text is generally reused in this document with little or no  
20 modification.

21 In the following sections, text in normal font describes required aspects of the  
22 DWARF format. Text in *italics* is explanatory or supplementary material, and not  
23 part of the format definition itself. The several appendices consist only of  
24 explanatory or supplementary material, and are not part of the formal definition.

### 25 **1.3 Objectives and Rationale**

26 DWARF has had a set of objectives since its inception which have guided the  
27 design and evolution of the debugging format. A discussion of these objectives  
28 and the rationale behind them may help with an understanding of the DWARF  
29 Debugging Format.

30 Although DWARF Version 1 was developed in the late 1980's as a format to  
31 support debugging C programs written for AT&T hardware running SVR4,  
32 DWARF Version 2 and later has evolved far beyond this origin. One difference  
33 between DWARF and other formats is that the latter are often specific to a  
34 particular language, architecture, and/or operating system.



### 1.3.1 Language Independence

DWARF is applicable to a broad range of existing procedural languages and is designed to be extensible to future languages. These languages may be considered to be "C-like" but the characteristics of C are not incorporated into DWARF Version 2 and later, unlike DWARF Version 1 and other debugging formats. DWARF abstracts concepts as much as possible so that the description can be used to describe a program in any language. As an example, the DWARF descriptions used to describe C functions, Pascal subroutines, and Fortran subprograms are all the same, with different attributes used to specify the differences between these similar programming language features.

On occasion, there is a feature which is specific to one particular language and which doesn't appear to have more general application. For these, DWARF has a description designed to meet the language requirements, although, to the extent possible, an effort is made to generalize the attribute. An example of this is the DW\_TAG\_condition debugging information entry, used to describe COBOL level 88 conditions, which is described in abstract terms rather than COBOL-specific terms. Conceivably, this TAG might be used with a different language which had similar functionality.

### 1.3.2 Architecture Independence

DWARF can be used with a wide range of processor architectures, whether byte or word oriented, with any word or byte size. DWARF can be used with Von Neumann architectures, using a single address space for both code and data; Harvard architectures, with separate code and data address spaces; and potentially for other architectures such as DSPs with their idiosyncratic memory organizations. DWARF can be used with common register-oriented architectures or with stack architectures.

DWARF assumes that memory has individual units (words or bytes) which have unique addresses which are ordered. (Identifying aliases is an implementation issue.)

### 1.3.3 Operating System Independence

DWARF is widely associated with SVR4 Unix and similar operating systems like BSD and Linux. DWARF fits well with the section organization of the ELF object file format. Nonetheless, DWARF attempts to be independent of either the OS or the object file format. There have been implementations of DWARF debugging data in COFF, Mach-O and other object file formats.

## Chapter 1. Introduction

1 DWARF assumes that any object file format will be able to distinguish the  
2 various DWARF data sections in some fashion, preferably by name.

3 DWARF makes a few assumptions about functionality provided by the  
4 underlying operating system. DWARF data sections can be read sequentially and  
5 independently. Each DWARF data section is a sequence of 8-bit bytes, numbered  
6 starting with zero. The presence of offsets from one DWARF data section into  
7 other data sections does not imply that the underlying OS must be able to  
8 position files randomly; a data section could be read sequentially and indexed  
9 using the offset.

### 10 **1.3.4 Compact Data Representation**

11 The DWARF description is designed to be a compact file-oriented representation.

12 There are several encodings which achieve this goal, such as the TAG and  
13 attribute abbreviations or the line number encoding. References from one section  
14 to another, especially to refer to strings, allow these sections to be compacted to  
15 eliminate duplicate data.

16 There are multiple schemes for eliminating duplicate data or reducing the size of  
17 the DWARF debug data associated with a given file. These include COMDAT,  
18 used to eliminate duplicate function or data definitions, the split DWARF object  
19 files which allow a consumer to find DWARF data in files other than the  
20 executable, or the type units, which allow similar type definitions from multiple  
21 compilations to be combined.

22 In most cases, it is anticipated that DWARF debug data will be read by a  
23 consumer (usually a debugger) and converted into a more efficiently accessed  
24 internal representation. For the most part, the DWARF data in a section is not the  
25 same as this internal representation.

### 26 **1.3.5 Efficient Processing**

27 DWARF is designed to be processed efficiently, so that a producer (a compiler)  
28 can generate the debug descriptions incrementally and a consumer can read only  
29 the descriptions which it needs at a given time. The data formats are designed to  
30 be efficiently interpreted by a consumer.

31 As mentioned, there is a tension between this objective and the preceding one. A  
32 DWARF data representation which resembles an internal data representation  
33 may lead to faster processing, but at the expense of larger data files. This may  
34 also constrain the possible implementations.

### 1.3.6 Implementation Independence

DWARF attempts to allow developers the greatest flexibility in designing implementations, without mandating any particular design decisions. Issues which can be described as quality-of-implementation are avoided.

### 1.3.7 Explicit Rather Than Implicit Description

DWARF describes the source to object translation explicitly rather than using common practice or convention as an implicit understanding between producer and consumer. For example, where other debugging formats assume that a debugger knows how to virtually unwind the stack, moving from one stack frame to the next using implicit knowledge about the architecture or operating system, DWARF makes this explicit in the Call Frame Information description.

### 1.3.8 Avoid Duplication of Information

DWARF has a goal of describing characteristics of a program once, rather than repeating the same information multiple times. The string sections can be compacted to eliminate duplicate strings, for example. Other compaction schemes or references between sections support this. Whether a particular implementation is effective at eliminating duplicate data, or even attempts to, is a quality-of-implementation issue.

### 1.3.9 Leverage Other Standards

Where another standard exists which describes how to interpret aspects of a program, DWARF defers to that standard rather than attempting to duplicate the description. For example, C++ has specific rules for deciding which function to call depending name, scope, argument types, and other factors. DWARF describes the functions and arguments, but doesn't attempt to describe how one would be selected by a consumer performing any particular operation.

### 1.3.10 Limited Dependence on Tools

DWARF data is designed so that it can be processed by commonly available assemblers, linkers, and other support programs, without requiring additional functionality specifically to support DWARF data. This may require the implementer to be careful that they do not generate DWARF data which cannot be processed by these programs. Conversely, an assembler which can generate LEB128 (Little-Endian Base 128) values may allow the compiler to generate more

1 compact descriptions, and a linker which understands the format of string  
2 sections can merge these sections. Whether or not an implementation includes  
3 these functions is a quality-of-implementation issue, not mandated by the  
4 DWARF specification.

### 5 **1.3.11 Separate Description From Implementation**

6 DWARF intends to describe the translation of a program from source to object,  
7 while neither mandating any particular design nor making any other design  
8 difficult. For example, DWARF describes how the arguments and local variables  
9 in a function are to be described, but doesn't specify how this data is collected or  
10 organized by a producer. Where a particular DWARF feature anticipates that it  
11 will be implemented in a certain fashion, informative text will suggest but not  
12 require this design.

### 13 **1.3.12 Permissive Rather Than Prescriptive**

14 The DWARF Standard specifies the meaning of DWARF descriptions. It does not  
15 specify in detail what a particular producer must generate for any source to  
16 object conversion. One producer may generate a more complete description than  
17 another, it may describe features in a different order (unless the standard  
18 explicitly requires a particular order), or it may use different abbreviations or  
19 compression methods. Similarly, DWARF does not specify exactly what a  
20 particular consumer should do with each part of the description, although we  
21 believe that the potential uses for each description should be evident.

22 DWARF is permissive, allowing different producers to generate different  
23 descriptions for the same source to object conversion, and permitting different  
24 consumers to provide more or less functionality or information to the user. This  
25 may result in debugging information being larger or smaller, compilers or  
26 debuggers which are faster or slower, and more or less functional. These are  
27 described as differences in quality-of-implementation.

28 Each producer conforming to the DWARF standard must follow the format and  
29 meaning as specified in the standard. As long as the DWARF description  
30 generated follows this specification, the producer is generating valid DWARF.  
31 For example, DWARF allows a producer to identify the end of a function  
32 prologue in the Line Information so that a debugger can stop at this location. A  
33 producer which does this is generating valid DWARF, as is another which  
34 doesn't. As another example, one producer may generate descriptions for  
35 variables which are moved from memory to a register in a certain range, while  
36 another may only describe the variable's location in memory. Both are valid

## Chapter 1. Introduction

1 DWARF descriptions, while a consumer using the former would be able to  
2 provide more accurate values for the variable while executing in that range than  
3 a consumer using the latter.

4 In this document, where the word “may” is used, the producer has the option to  
5 follow the description or not. Where the text says “may not”, this is prohibited.  
6 Where the text says “should”, this is advice about best practice, but is not a  
7 requirement.

### 8 **1.3.13 Vendor Extensibility**

9 This document does not attempt to cover all interesting languages or even to  
10 cover all of the possible debugging information needs for its primary target  
11 languages. Therefore, the document provides vendors a way to define their own  
12 debugging information tags, attributes, base type encodings, location operations,  
13 language names, calling conventions and call frame instructions by reserving a  
14 subset of the valid values for these constructs for vendor specific additions and  
15 defining related naming conventions. Vendors may also use debugging  
16 information entries and attributes defined here in new situations. Future  
17 versions of this document will not use names or values reserved for vendor  
18 specific additions. All names and values not reserved for vendor additions,  
19 however, are reserved for future versions of this document.

20 Where this specification provides a means for describing the source language,  
21 implementors are expected to adhere to that specification. For language features  
22 that are not supported, implementors may use existing attributes in novel ways  
23 or add vendor-defined attributes. Implementors who make extensions are  
24 strongly encouraged to design them to be compatible with this specification in  
25 the absence of those extensions.

1 The DWARF format is organized so that a consumer can skip over data which it  
2 does not recognize. This may allow a consumer to read and process files  
3 generated according to a later version of this standard or which contain vendor  
4 extensions, albeit possibly in a degraded manner.

## 5 **1.4 Changes from Version 5 to Version 6**

6 To be written...

## 7 **1.5 Changes from Version 4 to Version 5**

8 The following is a list of the major changes made to the DWARF Debugging  
9 Information Format since Version 4 was published. The list is not meant to be  
10 exhaustive.

- 11 • Eliminate the `.debug_types` section introduced in DWARF Version 4 and  
12 move its contents into the `.debug_info` section.
- 13 • Add support for collecting common DWARF information (debugging  
14 information entries and macro definitions) across multiple executable and  
15 shared files and keeping it in a single supplementary object file.
- 16 • Replace the line number program header format with a new format that  
17 provides the ability to use an MD5 hash to validate the source file version in  
18 use, allows pooling of directory and file name strings and makes provision  
19 for vendor-defined extensions. Also add a string section specific to the line  
20 number table (`.debug_line_str`) to properly support the common practice  
21 of stripping all DWARF sections except for line number information.
- 22 • Add a split object file and package representations to allow most DWARF  
23 information to be kept separate from an executable or shared image. This  
24 includes new sections `.debug_addr`, `.debug_str_offsets`,  
25 `.debug_abbrev.dwo`, `.debug_info.dwo`, `.debug_line.dwo`,  
26 `.debug_loclists.dwo`, `.debug_macro.dwo`, `.debug_str.dwo`,  
27 `.debug_str_offsets.dwo`, `.debug_cu_index` and `.debug_tu_index` together  
28 with new forms of attribute value for referencing these sections. This  
29 enhances DWARF support by reducing executable program size and by  
30 improving link times.
- 31 • Replace the `.debug_macinfo` macro information representation with with a  
32 `.debug_macro` representation that can potentially be much more compact.

## Chapter 1. Introduction

- 1 • Replace the `.debug_pubnames` and `.debug_pubtypes` sections with a single  
2 and more functional name index section, `.debug_names`.
- 3 • Replace the location list and range list sections (`.debug_loc` and  
4 `.debug_ranges`, respectively) with new sections (`.debug_loclists` and  
5 `.debug_rnglists`) and new representations that save space and processing  
6 time by eliminating most related object file relocations.
- 7 • Add a new debugging information entry (`DW_TAG_call_site`), related  
8 attributes and DWARF expression operators to describe call site  
9 information, including identification of tail calls and tail recursion.
- 10 • Add improved support for FORTRAN assumed rank arrays  
11 (`DW_TAG_generic_subrange`), dynamic rank arrays (`DW_AT_rank`) and  
12 co-arrays (`DW_TAG_coarray_type`).
- 13 • Add new operations that allow support for a DWARF expression stack  
14 containing typed values.
- 15 • Add improved support for the C++: auto return type, deleted member  
16 functions (`DW_AT_deleted`), as well as defaulted constructors and  
17 destructors (`DW_AT_defaulted`).
- 18 • Add a new attribute (`DW_AT_noreturn`), to identify a subprogram that  
19 does not return to its caller.
- 20 • Add language codes for C 2011, C++ 2003, C++ 2011, C++ 2014, Dylan,  
21 Fortran 2003, Fortran 2008, Go, Haskell, Julia, Modula 3, Ocaml, OpenCL  
22 C<sup>1</sup>, Rust and Swift.
- 23 • Numerous other more minor additions to improve functionality and  
24 performance.

25 DWARF Version 5 is compatible with DWARF Version 4 except as follows:

- 26 • The compilation unit header (in the `.debug_info` section) has a new  
27 `unit_type` field. In addition, the `debug_abbrev_offset` and `address_size`  
28 fields are reordered.
- 29 • New operand forms for attribute values are defined (`DW_FORM_addrx`,  
30 `DW_FORM_addrx1`, `DW_FORM_addrx2`, `DW_FORM_addrx3`,  
31 `DW_FORM_addrx4`, `DW_FORM_data16`, `DW_FORM_implicit_const`,  
32 `DW_FORM_line_strp`, `DW_FORM_loclistx`, `DW_FORM_rnglistx`,  
33 `DW_FORM_ref_sup4`, `DW_FORM_ref_sup8`, `DW_FORM_strp_sup`,

---

<sup>1</sup>called simply OpenCL in DWARF Version 5

## Chapter 1. Introduction

1 DW\_FORM\_strx, DW\_FORM\_strx1, DW\_FORM\_strx2, DW\_FORM\_strx3  
2 and DW\_FORM\_strx4.

3 *Because a pre-DWARF Version 5 consumer will not be able to interpret these even*  
4 *to ignore and skip over them, new forms must be considered incompatible additions.*

- 5 • The line number table header is substantially revised.
- 6 • The `.debug_loc` and `.debug_ranges` sections are replaced by new  
7 `.debug_loclists` and `.debug_rnglists` sections, respectively. These new  
8 sections have a new (and more efficient) list structure. Attributes that  
9 reference the predecessor sections must be interpreted differently to access  
10 the new sections. The new sections encode the same information as their  
11 predecessors, except that a new default location list entry is added.
- 12 • In a string type, the `DW_AT_byte_size` attribute is re-defined to always  
13 describe the size of the string type. (Previously `DW_AT_byte_size`  
14 described the size of the optional string length data field if the  
15 `DW_AT_string_length` attribute was also present.) In addition, the  
16 `DW_AT_string_length` attribute may now refer directly to an object that  
17 contains the length value.

18 While not strictly an incompatibility, the macro information representation is  
19 completely new; further, producers and consumers may optionally continue to  
20 support the older representation. While the two representations cannot both be  
21 used in the same compilation unit, they can co-exist in executable or shared  
22 images.

23 Similar comments apply to replacement of the `.debug_pubnames` and  
24 `.debug_pubtypes` sections with the new `.debug_names` section.

### 25 **1.6 Changes from Version 3 to Version 4**

26 The following is a list of the major changes made to the DWARF Debugging  
27 Information Format since Version 3 was published. The list is not meant to be  
28 exhaustive.

- 29 • Reformulate Section 2.6 (Location Descriptions) to better distinguish  
30 DWARF location descriptions, which compute the location where a value is  
31 found (such as an address in memory or a register name) from DWARF  
32 expressions, which compute a final value (such as an array bound).
- 33 • Add support for bundled instructions on machine architectures where  
34 instructions do not occupy a whole number of bytes.



## Chapter 1. Introduction

- 1 • Add a new attribute form for section offsets, `DW_FORM_sec_offset`, to  
2 replace the use of `DW_FORM_data4` and `DW_FORM_data8` for section  
3 offsets.
- 4 • Add an attribute, `DW_AT_main_subprogram`, to identify the main  
5 subprogram of a program.
- 6 • Define default array lower bound values for each supported language.
- 7 • Add a new technique using separate type units, type signatures and  
8 COMDAT sections to improve compression and duplicate elimination of  
9 DWARF information.
- 10 • Add support for new C++ language constructs, including rvalue references,  
11 generalized constant expressions, Unicode character types and template  
12 aliases.
- 13 • Clarify and generalize support for packed arrays and structures.
- 14 • Add new line number table support to facilitate profile based compiler  
15 optimization.
- 16 • Add additional support for template parameters in instantiations.
- 17 • Add support for strongly typed enumerations in languages (such as C++)  
18 that have two kinds of enumeration declarations.
- 19 • Add the option for the `DW_AT_high_pc` value of a program unit or scope  
20 to be specified as a constant offset relative to the corresponding  
21 `DW_AT_low_pc` value.

22 DWARF Version 4 is compatible with DWARF Version 3 except as follows:

- 23 • DWARF attributes that use any of the new forms of attribute value  
24 representation (for section offsets, flag compression, type signature  
25 references, and so on) cannot be read by DWARF Version 3 consumers  
26 because the consumer will not know how to skip over the unexpected form  
27 of data.
- 28 • DWARF frame and line number table sections include additional fields that  
29 affect the location and interpretation of other data in the section.

### 30 1.7 Changes from Version 2 to Version 3

31 The following is a list of the major differences between Version 2 and Version 3 of  
32 the DWARF Debugging Information Format. The list is not meant to be  
33 exhaustive.

## Chapter 1. Introduction

- 1 • Make provision for DWARF information files that are larger than 4 GBytes.
- 2 • Allow attributes to refer to debugging information entries in other shared
- 3 libraries.
- 4 • Add support for Fortran 90 modules as well as allocatable array and
- 5 pointer types.
- 6 • Add additional base types for C (as revised for 1999).
- 7 • Add support for Java and COBOL.
- 8 • Add namespace support for C++.
- 9 • Add an optional section for global type names (similar to the global section
- 10 for objects and functions).
- 11 • Adopt UTF-8 as the preferred representation of program name strings.
- 12 • Add improved support for optimized code (discontiguous scopes, end of
- 13 prologue determination, multiple section code generation).
- 14 • Improve the ability to eliminate duplicate DWARF information during
- 15 linking.

16 DWARF Version 3 is compatible with DWARF Version 2 except as follows:

- 17 • Certain very large values of the initial length fields that begin DWARF
- 18 sections as well as certain structures are reserved to act as escape codes for
- 19 future extension; one such extension is defined to increase the possible size
- 20 of DWARF descriptions (see [Section 7.4 on page 204](#)).
- 21 • References that use the attribute form DW\_FORM\_ref\_addr are specified to
- 22 be four bytes in the DWARF 32-bit format and eight bytes in the DWARF
- 23 64-bit format, while DWARF Version 2 specifies that such references have
- 24 the same size as an address on the target system (see [Sections 7.4 on](#)
- 25 [page 204](#) and [7.5.4 on page 218](#)).
- 26 • The return\_address\_register field in a Common Information Entry record
- 27 for call frame information is changed to unsigned LEB representation (see
- 28 [Section 6.4.1 on page 180](#)).

### 29 **1.8 Changes from Version 1 to Version 2**

30 DWARF Version 2 describes the second generation of debugging information  
31 based on the DWARF format. While DWARF Version 2 provides new debugging  
32 information not available in Version 1, the primary focus of the changes for

## Chapter 1. Introduction

1 Version 2 is the representation of the information, rather than the information  
2 content itself. The basic structure of the Version 2 format remains as in Version 1:  
3 the debugging information is represented as a series of debugging information  
4 entries, each containing one or more attributes (name/value pairs). The Version 2  
5 representation, however, is much more compact than the Version 1  
6 representation. In some cases, this greater density has been achieved at the  
7 expense of additional complexity or greater difficulty in producing and  
8 processing the DWARF information. The definers believe that the reduction in  
9 I/O and in memory paging should more than make up for any increase in  
10 processing time.

11 The representation of information changed from Version 1 to Version 2, so that  
12 Version 2 DWARF information is not binary compatible with Version 1  
13 information. To make it easier for consumers to support both Version 1 and  
14 Version 2 DWARF information, the Version 2 information has been moved to a  
15 different object file section, `.debug_info`.

16 *A summary of the major changes made in DWARF Version 2 compared to the DWARF*  
17 *Version 1 may be found in the DWARF Version 2 document.*

## Chapter 1. Introduction

*(empty page)*

# Chapter 2

## General Description

### 2.1 The Debugging Information Entry (DIE)

DWARF uses a series of debugging information entries (DIEs) to define a low-level representation of a source program. Each debugging information entry consists of an identifying tag and a series of attributes. An entry, or group of entries together, provide a description of a corresponding entity in the source program. The tag specifies the class to which an entry belongs and the attributes define the specific characteristics of the entry.

The set of tag names is listed in [Table 2.1 on the following page](#). The debugging information entries they identify are described in Chapters 3, 4 and 5.

*The debugging information entry descriptions in Chapters 3, 4 and 5 generally include mention of most, but not necessarily all, of the attributes that are normally or possibly used with the entry. Some attributes, whose applicability tends to be pervasive and invariant across many kinds of debugging information entries, are described in this section and not necessarily mentioned in all contexts where they may be appropriate. Examples include [DW\\_AT\\_artificial](#), the [declaration coordinates](#), and [DW\\_AT\\_description](#), among others.*

The debugging information entries are contained in the `.debug_info` and/or `.debug_info.dwo` sections of an object file.

Optionally, debugging information may be partitioned such that the majority of the debugging information can remain in individual object files without being processed by the linker. See [Section 7.3.2 on page 195](#) and [Appendix F on page 412](#) for details.

## Chapter 2. General Description

Table 2.1: Tag names

---

DW_TAG_access_declaration	DW_TAG_module
DW_TAG_array_type	DW_TAG_namelist
DW_TAG_atomic_type	DW_TAG_namelist_item
DW_TAG_base_type	DW_TAG_namespace
DW_TAG_call_site	DW_TAG_packed_type
DW_TAG_call_site_parameter	DW_TAG_partial_unit
DW_TAG_catch_block	DW_TAG_pointer_type
DW_TAG_class_type	DW_TAG_ptr_to_member_type
DW_TAG_coarray_type	DW_TAG_reference_type
DW_TAG_common_block	DW_TAG_restrict_type
DW_TAG_common_inclusion	DW_TAG_rvalue_reference_type
DW_TAG_compile_unit	DW_TAG_set_type
DW_TAG_condition	DW_TAG_shared_type
DW_TAG_const_type	DW_TAG_skeleton_unit
DW_TAG_constant	DW_TAG_string_type
DW_TAG_dwarf_procedure	DW_TAG_structure_type
DW_TAG_dynamic_type	DW_TAG_subprogram
DW_TAG_entry_point	DW_TAG_subrange_type
DW_TAG_enumeration_type	DW_TAG_subroutine_type
DW_TAG_enumerator	DW_TAG_template_alias
DW_TAG_file_type	DW_TAG_template_type_parameter
DW_TAG_formal_parameter	DW_TAG_template_value_parameter
DW_TAG_friend	DW_TAG_thrown_type
DW_TAG_generic_subrange	DW_TAG_try_block
DW_TAG_immutable_type	DW_TAG_typedef
DW_TAG_imported_declaration	DW_TAG_type_unit
DW_TAG_imported_module	DW_TAG_union_type
DW_TAG_imported_unit	DW_TAG_unspecified_parameters
DW_TAG_inheritance	DW_TAG_unspecified_type
DW_TAG_inlined_subroutine	DW_TAG_variable
DW_TAG_interface_type	DW_TAG_variant
DW_TAG_label	DW_TAG_variant_part
DW_TAG_lexical_block	DW_TAG_volatile_type
DW_TAG_member	DW_TAG_with_stmt

---

As a further option, debugging information entries and other debugging information that are the same in multiple executable or shared object files may be found in a separate supplementary object file that contains supplementary debug sections. See Section 7.3.6 on page 203 for further details.

## 2.2 Attribute Types

Each attribute value is characterized by an attribute name. No more than one attribute with a given name may appear in any debugging information entry. There are no limitations on the ordering of attributes within a debugging information entry.

The attributes are listed in Table 2.2 following.

Table 2.2: Attribute names

Attribute*	Usage
<a href="#">DW_AT_abstract_origin</a>	<a href="#">Inline instances of inline subprograms</a> <a href="#">Out-of-line instances of inline subprograms</a>
<a href="#">DW_AT_accessibility</a>	<a href="#">Access declaration (C++, Ada)</a> <a href="#">Accessibility of base or inherited class (C++)</a> <a href="#">Accessibility of data member or member function</a>
<a href="#">DW_AT_address_class</a>	<a href="#">Pointer or reference types</a> <a href="#">Subroutine or subroutine type</a>
<a href="#">DW_AT_addr_base</a>	<a href="#">Base offset for address table</a>
<a href="#">DW_AT_alignment</a>	<a href="#">Non-default alignment of type, subprogram or variable</a>
<a href="#">DW_AT_allocated</a>	<a href="#">Allocation status of types</a>
<a href="#">DW_AT_artificial</a>	<a href="#">Objects or types that are not actually declared in the source</a>
<a href="#">DW_AT_associated</a>	<a href="#">Association status of types</a>
<a href="#">DW_AT_base_types</a>	<a href="#">Primitive data types of compilation unit</a>
<a href="#">DW_AT_bias</a>	<a href="#">Integer bias added to an encoded value</a>
<a href="#">DW_AT_binary_scale</a>	<a href="#">Binary scale factor for fixed-point type</a>

*Continued on next page*

\*Links for attributes come to the left column of this table; links in the right column "fan-out" to one or more descriptions.

## Chapter 2. General Description

Attribute*	Identifies or Specifies
DW_AT_bit_size	Size of a base type in bits
DW_AT_bit_stride	Size of a data member in bits Array element stride (of array type) Subrange stride (dimension of array type) Enumeration stride (dimension of array type)
DW_AT_byte_size	Size of a data object or data type in bytes
DW_AT_byte_stride	Array element stride (of array type) Subrange stride (dimension of array type) Enumeration stride (dimension of array type)
DW_AT_call_all_calls	All tail and normal calls in a subprogram are described by call site entries
DW_AT_call_all_source_calls	All tail, normal and inlined calls in a subprogram are described by call site and inlined subprogram entries
DW_AT_call_all_tail_calls	All tail calls in a subprogram are described by call site entries
DW_AT_call_column	Column position of inlined subroutine call Column position of call site of non-inlined call
DW_AT_call_data_location	Address of the value pointed to by an argument passed in a call
DW_AT_call_data_value	Value pointed to by an argument passed in a call
DW_AT_call_file	File containing inlined subroutine call File containing call site of non-inlined call
DW_AT_call_line	Line number of inlined subroutine call Line containing call site of non-inlined call
DW_AT_call_origin	Subprogram called in a call
DW_AT_call_parameter	Parameter entry in a call
DW_AT_call_pc	Address of the call instruction in a call
DW_AT_call_return_pc	Return address from a call
DW_AT_call_tail_call	Call is a tail call

*Continued on next page*

\*Links for attributes come to the left column of this table; links in the right column "fan-out" to one or more descriptions.



## Chapter 2. General Description

Attribute*	Identifies or Specifies
<a href="#">DW_AT_call_target</a>	Address of called routine in a call
<a href="#">DW_AT_call_target_clobbered</a>	Address of called routine, which may be clobbered, in a call
<a href="#">DW_AT_call_value</a>	Argument value passed in a call
<a href="#">DW_AT_calling_convention</a>	Calling convention for subprograms Calling convention for types
<a href="#">DW_AT_common_reference</a>	Common block usage
<a href="#">DW_AT_comp_dir</a>	Compilation directory
<a href="#">DW_AT_const_expr</a>	Compile-time constant object Compile-time constant function
<a href="#">DW_AT_const_value</a>	Constant object Enumeration literal value Template value parameter
<a href="#">DW_AT_containing_type</a>	Containing type of pointer to member type
<a href="#">DW_AT_count</a>	Elements of subrange type
<a href="#">DW_AT_data_bit_offset</a>	Base type bit location Data member bit location
<a href="#">DW_AT_data_location</a>	Indirection to actual data
<a href="#">DW_AT_data_member_location</a>	Data member location Inherited member location
<a href="#">DW_AT_decimal_scale</a>	Decimal scale factor
<a href="#">DW_AT_decimal_sign</a>	Decimal sign representation
<a href="#">DW_AT_decl_column</a>	Column position of source declaration
<a href="#">DW_AT_decl_file</a>	File containing source declaration
<a href="#">DW_AT_decl_line</a>	Line number of source declaration
<a href="#">DW_AT_declaration</a>	Incomplete, non-defining, or separate entity declaration
<a href="#">DW_AT_defaulted</a>	Whether a member function has been declared as default
<a href="#">DW_AT_default_value</a>	Default value of parameter
<a href="#">DW_AT_deleted</a>	Whether a member has been declared as deleted

*Continued on next page*

\*Links for attributes come to the left column of this table; links in the right column "fan-out" to one or more descriptions.

## Chapter 2. General Description

Attribute*	Identifies or Specifies
<a href="#">DW_AT_description</a>	Artificial name or description
<a href="#">DW_AT_digit_count</a>	Digit count for packed decimal or numeric string type
<a href="#">DW_AT_discr</a>	Discriminant of variant part
<a href="#">DW_AT_discr_list</a>	List of discriminant values
<a href="#">DW_AT_discr_value</a>	Discriminant value
<a href="#">DW_AT_dwo_name</a>	Name of split DWARF object file
<a href="#">DW_AT_elemental</a>	Elemental property of a subroutine
<a href="#">DW_AT_encoding</a>	Encoding of base type
<a href="#">DW_AT_endianity</a>	Endianity of data
<a href="#">DW_AT_entry_pc</a>	Entry address of a scope (compilation unit, subprogram, and so on)
<a href="#">DW_AT_enum_class</a>	Type safe enumeration definition
<a href="#">DW_AT_explicit</a>	Explicit property of member function
<a href="#">DW_AT_export_symbols</a>	Export (inline) symbols of namespace Export symbols of a structure, union or class
<a href="#">DW_AT_extension</a>	Previous namespace extension or original namespace
<a href="#">DW_AT_external</a>	External subroutine External variable
<a href="#">DW_AT_frame_base</a>	Subroutine frame base address
<a href="#">DW_AT_friend</a>	Friend relationship
<a href="#">DW_AT_high_pc</a>	Contiguous range of code addresses
<a href="#">DW_AT_identifier_case</a>	Identifier case rule
<a href="#">DW_AT_import</a>	Imported declaration Imported unit Namespace alias Namespace using declaration Namespace using directive
<a href="#">DW_AT_inline</a>	Abstract instance Inlined subroutine
<a href="#">DW_AT_is_optional</a>	Optional parameter

*Continued on next page*

\*Links for attributes come to the left column of this table; links in the right column "fan-out" to one or more descriptions.

## Chapter 2. General Description

Attribute*	Identifies or Specifies
DW_AT_language_name	Programming language name
DW_AT_language_version	Programming language version
DW_AT_linkage_name	Object file linkage name of an entity
DW_AT_location	Data object location
DW_AT_loclists_base	Location lists base
DW_AT_low_pc	Code address or range of addresses
	Base address of scope
DW_AT_lower_bound	Lower bound of subrange
DW_AT_macros	Macro preprocessor information ( <i>#define, #undef, and so on in C, C++ and similar languages</i> )
DW_AT_main_subprogram	Main or starting subprogram Unit containing main or starting subprogram
DW_AT_mutable	Mutable property of member data
DW_AT_name	Name of declaration Path name of compilation source
DW_AT_namelist_item	Namelist item
DW_AT_noreturn	"no return" property of a subprogram
DW_AT_object_pointer	Object ( <i>this, self</i> ) pointer of member function
DW_AT_ordering	Array row/column ordering
DW_AT_picture_string	Picture string for numeric string type
DW_AT_priority	Module priority
DW_AT_producer	Compiler identification
DW_AT_prototyped	Subroutine prototype
DW_AT_pure	Pure property of a subroutine
DW_AT_ranges	Non-contiguous range of code addresses
DW_AT_rank	Dynamic number of array dimensions
DW_AT_recursive	Recursive property of a subroutine
DW_AT_reference	&-qualified non-static member function (C++)
DW_AT_return_addr	Subroutine return address save location

*Continued on next page*

\*Links for attributes come to the left column of this table; links in the right column "fan-out" to one or more descriptions.

## Chapter 2. General Description

Attribute*	Identifies or Specifies
<a href="#">DW_AT_rnglists_base</a>	Base offset for range lists
<a href="#">DW_AT_rvalue_reference</a>	&&-qualified non-static member function (C++)
<a href="#">DW_AT_scale_divisor</a>	Denominator of rational scale factor
<a href="#">DW_AT_scale_multiplier</a>	Numerator of rational scale factor
<a href="#">DW_AT_sibling</a>	Debugging information entry relationship
<a href="#">DW_AT_signature</a>	Type signature
<a href="#">DW_AT_small</a>	Scale factor for fixed-point type
<a href="#">DW_AT_specification</a>	Incomplete, non-defining, or separate declaration corresponding to a declaration
<a href="#">DW_AT_start_scope</a>	Reduced scope of declaration
<a href="#">DW_AT_static_link</a>	Location of uplevel frame
<a href="#">DW_AT_stmt_list</a>	Line number information for unit
<a href="#">DW_AT_string_length</a>	String length of string type
<a href="#">DW_AT_string_length_bit_size</a>	Size of string length of string type
<a href="#">DW_AT_string_length_byte_size</a>	Size of string length of string type
<a href="#">DW_AT_str_offsets<sup>1</sup></a>	String offsets information for unit
<a href="#">DW_AT_tensor</a>	Tensor (array) type
<a href="#">DW_AT_threads_scaled</a>	Array bound THREADS scale factor (UPC)
<a href="#">DW_AT_trampoline</a>	Target subroutine
<a href="#">DW_AT_type</a>	Type of call site
	Type of string type components
	Type of subroutine return
	Type of declaration
<a href="#">DW_AT_upper_bound</a>	Upper bound of subrange
<a href="#">DW_AT_use_location</a>	Member location for pointer to member type
<a href="#">DW_AT_use_UTF8</a>	Compilation unit uses UTF-8 strings
<a href="#">DW_AT_variable_parameter</a>	Non-constant parameter flag

*Continued on next page*

\*Links for attributes come to the left column of this table; links in the right column "fan-out" to one or more descriptions.

<sup>1</sup> DW\_FORM\_str\_offsets is new in DWARF Version 6. It replaces DW\_AT\_str\_offsets\_base which is deprecated.

## Chapter 2. General Description

Attribute*	Identifies or Specifies
<a href="#">DW_AT_virtuality</a>	<a href="#">Virtuality attribute</a>
<a href="#">DW_AT_visibility</a>	<a href="#">Visibility of declaration</a>
<a href="#">DW_AT_vtable_elem_location</a>	<a href="#">Virtual function vtable slot</a>

\*Links for attributes come to the left column of this table; links in the right column "fan-out" to one or more descriptions.

1 The permissible values for an attribute belong to one or more classes of attribute  
 2 value forms. Each form class may be represented in one or more ways. For  
 3 example, some attribute values consist of a single piece of constant data.  
 4 "Constant data" is the class of attribute value that those attributes may have.  
 5 There are several representations of constant data, including fixed length data of  
 6 one, two, four, eight or 16 bytes in size, and variable length data). The particular  
 7 representation for any given instance of an attribute is encoded along with the  
 8 attribute name as part of the information that guides the interpretation of a  
 9 debugging information entry.

10 Attribute value forms belong to one of the classes shown in Table 2.3 following.

Table 2.3: Classes of attribute value

Attribute Class	General Use and Encoding
<a href="#">address</a>	Refers to some location in the address space of the described program.
<a href="#">addrptr</a>	Specifies a location in the DWARF section that holds a series of machine address values. Certain attributes use one of these addresses by indexing relative to this location.
<a href="#">block</a>	An arbitrary number of uninterpreted bytes of data. The number of data bytes may be implicit from context or explicitly specified by an initial unsigned LEB128 value (see Section 7.6 on page 232) that precedes that number of data bytes.
<a href="#">constant</a>	One, two, four, eight or sixteen bytes of uninterpreted data, or data encoded in the variable length format known as LEB128 (see Section 7.6 on page 232).
<a href="#">exprval</a>	A DWARF expression yielding a value (see Section 2.5 on page 26). A leading unsigned ULEB128 value (see Section 7.6 on page 232) specifies the number of bytes in the expression.

*Continued on next page*

Attribute Class	General Use and Encoding
<a href="#">flag</a>	A small constant that indicates the presence or absence of an attribute.
<a href="#">lineptr</a>	Specifies a location in the DWARF section that holds line number information.
<a href="#">locdesc</a>	A DWARF location description (see Section 2.6 on page 38). A leading unsigned ULEB128 value (see Section 7.6 on page 232) specifies the number of bytes in the location description.
<a href="#">loclist, loclistsptr</a>	Specifies a location in the DWARF section that holds location lists, which describe objects whose location can change during their lifetime.
<a href="#">macptr</a>	Specifies a location in the DWARF section that holds macro definition information.
<a href="#">reference</a>	Refers to one of the debugging information entries that describe the program. There are four types of reference. The first is an offset relative to the beginning of the compilation unit in which the reference occurs and must refer to an entry within that same compilation unit. The second type of reference is the offset of a debugging information entry in any compilation unit, including one different from the unit containing the reference. The third type of reference is an indirect reference to a type definition using an 8-byte signature for that type. The fourth type of reference is a reference from within the <code>.debug_info</code> section of the executable or shared object file to a debugging information entry in the <code>.debug_info</code> section of a supplementary object file.
<a href="#">rnglist, rnglistsptr</a>	Specifies a location in the DWARF section that holds non-contiguous address ranges.
<a href="#">string</a>	A null-terminated sequence of zero or more (non-null) bytes. Data in this class are generally printable strings. Strings may be represented directly in the debugging information entry or as an offset in a separate string table.

*Continued on next page*

Attribute Class	General Use and Encoding
<a href="#">stroffsetsptr</a>	Specifies a location in the DWARF section that holds a series of offsets into the DWARF section that holds strings. Certain attributes use one of these offsets by indexing relative to this location. The resulting offset is then used to index into the DWARF string section.

## 2.3 Relationship of Debugging Information Entries

*A variety of needs can be met by permitting a single debugging information entry to “own” an arbitrary number of other debugging entries and by permitting the same debugging information entry to be one of many owned by another debugging information entry. This makes it possible, for example, to describe the static [block](#) structure within a source file, to show the members of a structure, union, or class, and to associate declarations with source files or source files with shared object files.*

The ownership relationship of debugging information entries is achieved naturally because the debugging information is represented as a tree. The nodes of the tree are the debugging information entries themselves. The child entries of any node are exactly those debugging information entries owned by that node.

*While the ownership relation of the debugging information entries is represented as a tree, other relations among the entries exist, for example, a reference from an entry representing a variable to another entry representing the type of that variable. If all such relations are taken into account, the debugging entries form a graph, not a tree.*

The tree itself is represented by flattening it in prefix order. Each debugging information entry is defined either to have child entries or not to have child entries (see Section [7.5.3 on page 213](#)). If an entry is defined not to have children, the next physically succeeding entry is a sibling. If an entry is defined to have children, the next physically succeeding entry is its first child. Additional children are represented as siblings of the first child. A chain of sibling entries is terminated by a null entry.

In cases where a producer of debugging information feels that it will be important for consumers of that information to quickly scan chains of sibling entries, while ignoring the children of individual siblings, that producer may attach a [DW\\_AT\\_sibling](#) attribute to any debugging information entry. The value of this attribute is a reference to the sibling entry of the entry to which the attribute is attached.



## 2.4 Target Addresses

Addresses, bytes and bits in DWARF use the numbering and direction conventions that are appropriate to the current language on the target system.

Many places in this document refer to the size of an address on the target architecture (or equivalently, target machine) to which a DWARF description applies. For processors which can be configured to have different address sizes or different instruction sets, the intent is to refer to the configuration which is either the default for that processor or which is specified by the object file or executable file which contains the DWARF information.

*For example, if a particular target architecture supports both 32-bit and 64-bit addresses, the compiler will generate an object file which specifies that it contains executable code generated for one or the other of these address sizes. In that case, the DWARF debugging information contained in this object file will use the same address size.*

### 2.4.1 Reserved Target Address for Non-Existent Entity

The target address consisting of the largest representable address value (for example, 0xffffffff for a 32-bit address) is reserved to indicate that there is no entity designated by that address.

*In some cases a producer may emit machine code or allocate storage for an entity, but a linker or other subsequent processing step may remove that entity. In that case, rather than be required to rewrite the DWARF description to eliminate the relevant DWARF construct that contains the address of that entity, the processing step may simply update the address value to the reserved value.*

## 2.5 DWARF Expressions

DWARF expressions describe how to compute a value or specify a location. They are expressed in terms of DWARF operations that operate on a stack of values.

A DWARF expression is encoded as a stream of operations, each consisting of an opcode followed by zero or more literal operands. The number of operands is implied by the opcode.

In addition to the general operations that are defined here, operations that are specific to location descriptions are defined in [Section 2.6 on page 38](#).



## 2.5.1 General Operations

Each general operation represents a postfix operation on a simple stack machine. Each element of the stack has a type and a value, and can represent a value of any supported base type of the target machine. Instead of a base type, elements can have a **generic type**, which is an integral type that has the size of an address on the target machine and unspecified signedness. The value on the top of the stack after “executing” the DWARF expression is taken to be the result (the address of the object, the value of the array bound, the length of a dynamic string, the desired value itself, and so on).

*The **generic type** is the same as the unspecified type used for stack operations defined in DWARF Version 4 and before.*

### 2.5.1.1 Literal Encodings

The following operations all push a value onto the DWARF stack. Operations other than **DW\_OP\_const\_type** push a value with the **generic type**, and if the value of a constant in one of these operations is larger than can be stored in a single stack element, the value is truncated to the element size and the low-order bits are pushed on the stack.

1. **DW\_OP\_lit0, DW\_OP\_lit1, ..., DW\_OP\_lit31**  
The **DW\_OP\_lit<n>** operations encode the unsigned literal values from 0 through 31, inclusive.
2. **DW\_OP\_addr**  
The **DW\_OP\_addr** operation has a single operand that encodes a machine address and whose size is the size of an address on the target machine.
3. **DW\_OP\_const1u, DW\_OP\_const2u, DW\_OP\_const4u, DW\_OP\_const8u**  
The single operand of a **DW\_OP\_const<n>u** operation provides a 1, 2, 4, or 8-byte unsigned integer constant, respectively.
4. **DW\_OP\_const1s, DW\_OP\_const2s, DW\_OP\_const4s, DW\_OP\_const8s**  
The single operand of a **DW\_OP\_const<n>s** operation provides a 1, 2, 4, or 8-byte signed integer constant, respectively.
5. **DW\_OP\_constu**  
The single operand of the **DW\_OP\_constu** operation provides an unsigned LEB128 integer constant.
6. **DW\_OP\_consts**  
The single operand of the **DW\_OP\_consts** operation provides a signed LEB128 integer constant.

1       7. **DW\_OP\_addrx**

2       The DW\_OP\_addrx operation has a single operand that encodes an unsigned  
3       LEB128 value, which is a zero-based index into the `.debug_addr` section,  
4       where a machine address is stored. This index is relative to the value of the  
5       [DW\\_AT\\_addr\\_base](#) attribute of the associated compilation unit.

6       8. **DW\_OP\_constx**

7       The DW\_OP\_constx operation has a single operand that encodes an unsigned  
8       LEB128 value, which is a zero-based index into the `.debug_addr` section,  
9       where a constant, the size of a machine address, is stored. This index is  
10       relative to the value of the [DW\\_AT\\_addr\\_base](#) attribute of the associated  
11       compilation unit.

12       *The DW\_OP\_constx operation is provided for constants that require link-time  
13       relocation but should not be interpreted by the consumer as a relocatable address (for  
14       example, offsets to thread-local storage).*

15       9. **DW\_OP\_const\_type**

16       The DW\_OP\_const\_type operation takes three operands. The first operand is  
17       an unsigned LEB128 integer that represents the offset of a debugging  
18       information entry in the current compilation unit, which must be a  
19       [DW\\_TAG\\_base\\_type](#) entry that provides the type of the constant provided.  
20       The second operand is 1-byte unsigned integer that specifies the size of the  
21       constant value, which is the same as the size of the base type referenced by  
22       the first operand. The third operand is a sequence of bytes of the given size  
23       that is interpreted as a value of the referenced type.

24       *While the size of the constant can be inferred from the base type definition, it is  
25       encoded explicitly into the operation so that the operation can be parsed easily  
26       without reference to the `.debug_info` section.*

27       **2.5.1.2 Register Values**

28       The following operations push a value onto the stack that is either part or all of  
29       the contents of a register or the result of adding the contents of a register to a  
30       given signed offset. [DW\\_OP\\_regval\\_type](#) pushes the contents of a register  
31       together with the given base type. [DW\\_OP\\_regval\\_bits](#) pushes the partial  
32       contents of a register together with the generic type. The other operations push  
33       the result of adding the contents of a register to a given signed offset together  
34       with the [generic type](#).

## Chapter 2. General Description

### 1. **DW\_OP\_fbreg**

The DW\_OP\_fbreg operation provides a signed LEB128 offset from the address specified by the location description in the [DW\\_AT\\_frame\\_base](#) attribute of the current function.

*This is typically a stack pointer register plus or minus some offset.*

### 2. **DW\_OP\_breg0, DW\_OP\_breg1, ..., DW\_OP\_breg31**

The single operand of the [DW\\_OP\\_breg<n>](#) operations provides a signed LEB128 offset from the contents of the specified register.

### 3. **DW\_OP\_bregx**

The DW\_OP\_bregx operation provides the sum of two values specified by its two operands. The first operand is a register number which is specified by an unsigned LEB128 number. The second operand is a signed LEB128 offset.

### 4. **DW\_OP\_regval\_type**

The DW\_OP\_regval\_type operation pushes the contents of a given register interpreted as a value of a given type. The first operand is an unsigned LEB128 number, which identifies a register whose contents is to be pushed onto the stack. The second operand is an unsigned LEB128 number that represents the offset of a debugging information entry in the current compilation unit, which must be a [DW\\_TAG\\_base\\_type](#) entry that provides the type of the value contained in the specified register.

### 5. **DW\_OP\_regval\_bits**

The DW\_OP\_regval\_bits operation takes a single unsigned LEB128 integer operand, which gives the number of bits to read. This number must be smaller or equal to the bit size of the generic type. It pops the top two stack elements and interprets the top element as an unsigned bit offset from the least significant bit end and the other as a register number identifying the register from which to extract the value. If the extracted value is smaller than the size of the generic type, it is zero extended.

### 2.5.1.3 Stack Operations

The following operations manipulate the DWARF stack. Operations that index the stack assume that the top of the stack (most recently added entry) has index 0.

Each entry on the stack has an associated type.

#### 1. **DW\_OP\_dup**

The DW\_OP\_dup operation duplicates the value (including its type identifier) at the top of the stack.

## Chapter 2. General Description

- 1     2. **DW\_OP\_drop**  
2         The DW\_OP\_drop operation pops the value (including its type identifier) at  
3         the top of the stack.
- 4     3. **DW\_OP\_pick**  
5         The single operand of the DW\_OP\_pick operation provides a 1-byte index. A  
6         copy of the stack entry (including its type identifier) with the specified index  
7         (0 through 255, inclusive) is pushed onto the stack.
- 8     4. **DW\_OP\_over**  
9         The DW\_OP\_over operation duplicates the entry currently second in the  
10         stack at the top of the stack. This is equivalent to a DW\_OP\_pick operation,  
11         with index 1.
- 12    5. **DW\_OP\_swap**  
13         The DW\_OP\_swap operation swaps the top two stack entries. The entry at  
14         the top of the stack (including its type identifier) becomes the second stack  
15         entry, and the second entry (including its type identifier) becomes the top of  
16         the stack.
- 17    6. **DW\_OP\_rot**  
18         The DW\_OP\_rot operation rotates the first three stack entries. The entry at the  
19         top of the stack (including its type identifier) becomes the third stack entry,  
20         the second entry (including its type identifier) becomes the top of the stack,  
21         and the third entry (including its type identifier) becomes the second entry.
- 22    7. **DW\_OP\_deref**  
23         The DW\_OP\_deref operation pops the top stack entry and treats it as an  
24         address. The popped value must have an integral type. The value retrieved  
25         from that address is pushed, and has the [generic type](#). The size of the data  
26         retrieved from the dereferenced address is the size of an address on the target  
27         machine.
- 28    8. **DW\_OP\_deref\_size**  
29         The DW\_OP\_deref\_size operation behaves like the [DW\\_OP\\_deref](#) operation:  
30         it pops the top stack entry and treats it as an address. The popped value must  
31         have an integral type. The value retrieved from that address is pushed, and  
32         has the [generic type](#). In the DW\_OP\_deref\_size operation, however, the size  
33         in bytes of the data retrieved from the dereferenced address is specified by  
34         the single operand. This operand is a 1-byte unsigned integral constant  
35         whose value may not be larger than the size of the [generic type](#). The data  
36         retrieved is zero extended to the size of an address on the target machine  
37         before being pushed onto the expression stack.

### 9. **DW\_OP\_deref\_type**

The DW\_OP\_deref\_type operation behaves like the DW\_OP\_deref\_size operation: it pops the top stack entry and treats it as an address. The popped value must have an integral type. The value retrieved from that address is pushed together with a type identifier. In the DW\_OP\_deref\_type operation, the size in bytes of the data retrieved from the dereferenced address is specified by the first operand. This operand is a 1-byte unsigned integral constant whose value which is the same as the size of the base type referenced by the second operand. The second operand is an unsigned LEB128 integer that represents the offset of a debugging information entry in the current compilation unit, which must be a DW\_TAG\_base\_type entry that provides the type of the data pushed.

*While the size of the pushed value could be inferred from the base type definition, it is encoded explicitly into the operation so that the operation can be parsed easily without reference to the .debug\_info section.*

### 10. **DW\_OP\_xderef**

The DW\_OP\_xderef operation provides an extended dereference mechanism. The entry at the top of the stack is treated as an address. The second stack entry is treated as an “address space identifier” for those architectures that support multiple address spaces. Both of these entries must have integral type identifiers. The top two stack elements are popped, and a data item is retrieved through an implementation-defined address calculation and pushed as the new stack top together with the generic type identifier. The size of the data retrieved from the dereferenced address is the size of the generic type.

### 11. **DW\_OP\_xderef\_size**

The DW\_OP\_xderef\_size operation behaves like the DW\_OP\_xderef operation. The entry at the top of the stack is treated as an address. The second stack entry is treated as an “address space identifier” for those architectures that support multiple address spaces. Both of these entries must have integral type identifiers. The top two stack elements are popped, and a data item is retrieved through an implementation-defined address calculation and pushed as the new stack top. In the DW\_OP\_xderef\_size operation, however, the size in bytes of the data retrieved from the dereferenced address is specified by the single operand. This operand is a 1-byte unsigned integral constant whose value may not be larger than the size of an address on the target machine. The data retrieved is zero extended to the size of an address on the target machine before being pushed onto the expression stack together with the generic type identifier.

12. **DW\_OP\_xderef\_type**

The DW\_OP\_xderef\_type operation behaves like the DW\_OP\_xderef\_size operation: it pops the top two stack entries, treats them as an address and an address space identifier, and pushes the value retrieved. In the DW\_OP\_xderef\_type operation, the size in bytes of the data retrieved from the dereferenced address is specified by the first operand. This operand is a 1-byte unsigned integral constant whose value value which is the same as the size of the base type referenced by the second operand. The second operand is an unsigned LEB128 integer that represents the offset of a debugging information entry in the current compilation unit, which must be a DW\_TAG\_base\_type entry that provides the type of the data pushed.

13. **DW\_OP\_push\_object\_address**

The DW\_OP\_push\_object\_address operation pushes the address of the object currently being evaluated as part of evaluation of a user presented expression. This object may correspond to an independent variable described by its own debugging information entry or it may be a component of an array, structure, or class whose address has been dynamically determined by an earlier step during user expression evaluation.

*This operator provides explicit functionality (especially for arrays involving descriptors) that is analogous to the implicit push of the base address of a structure prior to evaluation of a DW\_AT\_data\_member\_location to access a data member of a structure. For an example, see Appendix D.2 on page 304.*

14. **DW\_OP\_form\_tls\_address**

The DW\_OP\_form\_tls\_address operation pops a value from the stack, which must have an integral type identifier, translates this value into an address in the thread-local storage for a thread, and pushes the address onto the stack together with the generic type identifier. The meaning of the value on the top of the stack prior to this operation is defined by the run-time environment. If the run-time environment supports multiple thread-local storage blocks for a single thread, then the block corresponding to the executable or shared library containing this DWARF expression is used.

*Some implementations of C, C++, Fortran, and other languages, support a thread-local storage class. Variables with this storage class have distinct values and addresses in distinct threads, much as automatic variables have distinct values and addresses in each function invocation. Typically, there is a single block of storage containing all thread-local variables declared in the main executable, and a separate block for the variables declared in each shared library. Each thread-local variable can then be accessed in its block using an identifier. This identifier is typically an offset into the block and pushed onto the DWARF stack by one of the*



## Chapter 2. General Description

*DW\_OP\_const<n><x> operations prior to the DW\_OP\_form\_tls\_address operation. Computing the address of the appropriate block can be complex (in some cases, the compiler emits a function call to do it), and difficult to describe using ordinary DWARF location descriptions. Instead of forcing complex thread-local storage calculations into the DWARF expressions, the DW\_OP\_form\_tls\_address allows the consumer to perform the computation based on the run-time environment.*

### 15. **DW\_OP\_call\_frame\_cfa**

The DW\_OP\_call\_frame\_cfa operation pushes the value of the CFA, obtained from the Call Frame Information (see Section 6.4 on page 178).

*Although the value of DW\_AT\_frame\_base can be computed using other DWARF expression operators, in some cases this would require an extensive location list because the values of the registers used in computing the CFA change during a subroutine. If the Call Frame Information is present, then it already encodes such changes, and it is space efficient to reference that.*

*Examples illustrating many of these stack operations are found in Appendix D.1.2 on page 300.*

### 2.5.1.4 Arithmetic and Logical Operations

The following provide arithmetic and logical operations. Operands of an operation with two operands must have the same type, either the same base type or the **generic type**. The result of the operation which is pushed back has the same type as the type of the operand(s).

If the type of the operands is the **generic type**, except as otherwise specified, the arithmetic operations perform addressing arithmetic, that is, unsigned arithmetic that is performed modulo one plus the largest representable address.

Operations other than DW\_OP\_abs, DW\_OP\_div, DW\_OP\_minus, DW\_OP\_mul, DW\_OP\_neg and DW\_OP\_plus require integral types of the operand (either integral base type or the **generic type**). Operations do not cause an exception on overflow.

#### 1. **DW\_OP\_abs**

The DW\_OP\_abs operation pops the top stack entry, interprets it as a signed value and pushes its absolute value. If the absolute value cannot be represented, the result is undefined.

#### 2. **DW\_OP\_and**

The DW\_OP\_and operation pops the top two stack values, performs a bitwise and operation on the two, and pushes the result.

## Chapter 2. General Description

- 1     3. **DW\_OP\_div**  
2       The DW\_OP\_div operation pops the top two stack values, divides the former  
3       second entry by the former top of the stack using signed division, and pushes  
4       the result.
- 5     4. **DW\_OP\_minus**  
6       The DW\_OP\_minus operation pops the top two stack values, subtracts the  
7       former top of the stack from the former second entry, and pushes the result.
- 8     5. **DW\_OP\_mod**  
9       The DW\_OP\_mod operation pops the top two stack values and pushes the  
10      result of the calculation: former second stack entry modulo the former top of  
11      the stack.
- 12    6. **DW\_OP\_mul**  
13      The DW\_OP\_mul operation pops the top two stack entries, multiplies them  
14      together, and pushes the result.
- 15    7. **DW\_OP\_neg**  
16      The DW\_OP\_neg operation pops the top stack entry, interprets it as a signed  
17      value and pushes its negation. If the negation cannot be represented, the  
18      result is undefined.
- 19    8. **DW\_OP\_not**  
20      The DW\_OP\_not operation pops the top stack entry, and pushes its bitwise  
21      complement.
- 22    9. **DW\_OP\_or**  
23      The DW\_OP\_or operation pops the top two stack entries, performs a bitwise  
24      or operation on the two, and pushes the result.
- 25    10. **DW\_OP\_plus**  
26       The DW\_OP\_plus operation pops the top two stack entries, adds them  
27       together, and pushes the result.
- 28    11. **DW\_OP\_plus\_uconst**  
29       The DW\_OP\_plus\_uconst operation pops the top stack entry, adds it to the  
30       unsigned LEB128 constant operand interpreted as the same type as the  
31       operand popped from the top of the stack and pushes the result.  
32       *This operation is supplied specifically to be able to encode more field offsets in two*  
33       *bytes than can be done with “DW\_OP\_lit<n> DW\_OP\_plus.”*



1 12. **DW\_OP\_shl**

2 The DW\_OP\_shl operation pops the top two stack entries, shifts the former  
3 second entry left (filling with zero bits) by the number of bits specified by the  
4 former top of the stack, and pushes the result.

5 13. **DW\_OP\_shr**

6 The DW\_OP\_shr operation pops the top two stack entries, shifts the former  
7 second entry right logically (filling with zero bits) by the number of bits  
8 specified by the former top of the stack, and pushes the result.

9 14. **DW\_OP\_shra**

10 The DW\_OP\_shra operation pops the top two stack entries, shifts the former  
11 second entry right arithmetically (divide the magnitude by 2, keep the same  
12 sign for the result) by the number of bits specified by the former top of the  
13 stack, and pushes the result.

14 15. **DW\_OP\_xor**

15 The DW\_OP\_xor operation pops the top two stack entries, performs a bitwise  
16 exclusive-or operation on the two, and pushes the result.

17 **2.5.1.5 Control Flow Operations**

18 The following operations provide simple control of the flow of a DWARF  
19 expression.

20 1. **DW\_OP\_le, DW\_OP\_ge, DW\_OP\_eq, DW\_OP\_lt, DW\_OP\_gt, DW\_OP\_ne**

21 The six relational operators each:

- 22 • pop the top two stack values, which have the same type, either the same  
23 base type or the [generic type](#),
- 24 • compare the operands:  
25 < former second entry >< relational operator >< former top entry >
- 26 • push the constant value 1 onto the stack if the result of the operation is  
27 true or the constant value 0 if the result of the operation is false. The  
28 pushed value has the [generic type](#).

29 If the operands have the [generic type](#), the comparisons are performed as  
30 signed operations.

### 2. **DW\_OP\_skip**

DW\_OP\_skip is an unconditional branch. Its single operand is a 2-byte signed integer constant. The 2-byte constant is the number of bytes of the DWARF expression to skip forward or backward from the current operation, beginning after the 2-byte constant.

### 3. **DW\_OP\_bra**

DW\_OP\_bra is a conditional branch. Its single operand is a 2-byte signed integer constant. This operation pops the top of stack. If the value popped is not the constant 0, the 2-byte constant operand is the number of bytes of the DWARF expression to skip forward or backward from the current operation, beginning after the 2-byte constant.

### 4. **DW\_OP\_call2, DW\_OP\_call4, DW\_OP\_call\_ref**

DW\_OP\_call2, DW\_OP\_call4, and DW\_OP\_call\_ref perform DWARF procedure calls during evaluation of a DWARF expression or location description. For DW\_OP\_call2 and DW\_OP\_call4, the operand is the 2- or 4-byte unsigned offset, respectively, of a debugging information entry in the current compilation unit. The DW\_OP\_call\_ref operator has a single operand. In the [32-bit DWARF format](#), the operand is a 4-byte unsigned value; in the [64-bit DWARF format](#), it is an 8-byte unsigned value (see [Section 7.4 on page 204](#)). The operand is used as the offset of a debugging information entry in the `.debug_info` section of the current executable or shared object file.

*Operand interpretation of DW\_OP\_call2, DW\_OP\_call4 and DW\_OP\_call\_ref is exactly like that for DW\_FORM\_ref2, DW\_FORM\_ref4 and DW\_FORM\_ref\_addr, respectively (see [Section 7.5.4 on page 218](#)).*

These operations transfer control of DWARF expression evaluation to the [DW\\_AT\\_location](#) attribute of the referenced debugging information entry. If there is no such attribute, then there is no effect. Execution of the DWARF expression of a [DW\\_AT\\_location](#) attribute may add to and/or remove from values on the stack. Execution returns to the point following the call when the end of the attribute is reached. Values on the stack at the time of the call may be used as parameters by the called expression and values left on the stack by the called expression may be used as return values by prior agreement between the calling and called expressions.

### 2.5.1.6 Type Conversions

The following operations provide for explicit type conversion.

#### 1. **DW\_OP\_convert**

The DW\_OP\_convert operation pops the top stack entry, converts it to a different type, then pushes the result. It takes one operand, which is an unsigned LEB128 integer that represents the offset of a debugging information entry in the current compilation unit, or value 0 which represents the [generic type](#). If the operand is non-zero, the referenced entry must be a [DW\\_TAG\\_base\\_type](#) entry that provides the type to which the value is converted.

#### 2. **DW\_OP\_reinterpret**

The DW\_OP\_reinterpret operation pops the top stack entry, reinterprets the bits in its value as a value of a different type, then pushes the result. It takes one operand, which is an unsigned LEB128 integer that represents the offset of a debugging information entry in the current compilation unit, or value 0 which represents the [generic type](#). If the operand is non-zero, the referenced entry must be a [DW\\_TAG\\_base\\_type](#) entry that provides the type to which the value is converted. The type of the operand and result type must have the same size in bits.

### 2.5.1.7 Special Operations

There are these special operations currently defined:

#### 1. **DW\_OP\_nop**

The DW\_OP\_nop operation is a place holder. It has no effect on the location stack or any of its values.

#### 2. **DW\_OP\_entry\_value**

The DW\_OP\_entry\_value operation pushes the value that an expression would have had, or a register location would have held, upon entering the current subprogram. It has two operands: an unsigned LEB128 length, followed by a block containing a DWARF expression or a register location description (see [Section 2.6.1.1.3 on page 39](#)). The length operand specifies the length in bytes of the block. If the block contains a DWARF expression, the DWARF expression is evaluated as if it had been evaluated upon entering the current subprogram. The DWARF expression assumes no values are present on the DWARF stack initially and results in exactly one value being pushed on the DWARF stack when completed. If the block contains a register location description, DW\_OP\_entry\_value pushes the value that register held upon entering the current subprogram.

1 [DW\\_OP\\_push\\_object\\_address](#) is not meaningful inside of this DWARF  
2 operation.

3 *The register location description provides a more compact form for the case where the*  
4 *value was in a register on entry to the subprogram.*

5 *The values needed to evaluate [DW\\_OP\\_entry\\_value](#) could be obtained in several*  
6 *ways. The consumer could suspend execution on entry to the subprogram, record*  
7 *values needed by [DW\\_OP\\_entry\\_value](#) expressions within the subprogram, and then*  
8 *continue; when evaluating [DW\\_OP\\_entry\\_value](#), the consumer would use these*  
9 *recorded values rather than the current values. Or, when evaluating*  
10 *[DW\\_OP\\_entry\\_value](#), the consumer could virtually unwind using the Call Frame*  
11 *Information (see Section 6.4 on page 178) to recover register values that might have*  
12 *been clobbered since the subprogram entry point.*

13 3. **DW\_OP\_user\_extended**

14 The [DW\\_OP\\_user\\_extended](#) opcode encodes a vendor extension operation. It  
15 has at least one operand: a ULEB128 constant identifying a vendor extension  
16 operation. The remaining operands are defined by the vendor extension. The  
17 vendor extension opcode 0 is reserved and cannot be used by any vendor  
18 extension.

19 *The [DW\\_OP\\_user\\_extended](#) encoding space can be understood to supplement the*  
20 *space defined by [DW\\_OP\\_lo\\_user](#) and [DW\\_OP\\_hi\\_user](#) that is allocated by the*  
21 *standard for the same purpose.*

22 **2.6 Location Descriptions**

23 *Debugging information must provide consumers a way to find the location of program*  
24 *variables, determine the bounds of dynamic arrays and strings, and possibly to find the*  
25 *base address of a subroutine's stack frame or the return address of a subroutine.*

26 *Furthermore, to meet the needs of recent computer architectures and optimization*  
27 *techniques, debugging information must be able to describe the location of an object*  
28 *whose location changes over the object's lifetime.*

29 Information about the location of program objects is provided by location  
30 descriptions. Location descriptions can be either of two forms:

- 31 1. *Single location descriptions*, which are a language independent representation  
32 of addressing rules of arbitrary complexity built from DWARF expressions  
33 (See Section 2.5 on page 26) and/or other DWARF operations specific to  
34 describing locations. They are sufficient for describing the location of any  
35 object as long as its lifetime is either static or the same as the [lexical block](#) that  
36 owns it, and it does not move during its lifetime.

2. *Location lists*, which are used to describe objects that have a limited lifetime or change their location during their lifetime. Location lists are described in Section 2.6.2 on page 43 below.

Location descriptions are distinguished in a context sensitive manner. As the value of an attribute, a single location description is encoded using class `locdesc` and a location list is encoded using class `loclist` (which serves as an index into a separate section containing location lists).

## 2.6.1 Single Location Descriptions

A single location description is either:

1. A simple location description, representing an object which exists in one contiguous piece at the given location, or
2. A composite location description consisting of one or more simple location descriptions, each of which is followed by one composition operation. Each simple location description describes the location of one piece of the object; each composition operation describes which part of the object is located there. Each simple location description that is a DWARF expression is evaluated independently of any others.

### 2.6.1.1 Simple Location Descriptions

A simple location description represents one contiguous piece or all of an object or value.

#### 2.6.1.1.1 Empty Location Descriptions

An empty location description consists of a DWARF expression containing no operations. It represents a piece or all of an object that is present in the source but not in the object code (perhaps due to optimization).

#### 2.6.1.1.2 Memory Location Descriptions

A memory location description consists of a non-empty DWARF expression (see Section 2.5 on page 26), whose value is the address of a piece or all of an object or other entity in memory.

#### 2.6.1.1.3 Register Location Descriptions

A register location description consists of a register name operation, which represents a piece or all of an object located in a given register.

## Chapter 2. General Description

1 *Register location descriptions describe an object (or a piece of an object) that resides in a*  
2 *register, while the opcodes listed in Section 2.5.1.2 on page 28 are used to describe an*  
3 *object (or a piece of an object) that is located in memory at an address that is contained in*  
4 *a register (possibly offset by some constant). A register location description must stand*  
5 *alone as the entire description of an object or a piece of an object.*

6 The following DWARF operations can be used to specify a register location.

7 *Note that the register number represents a DWARF specific mapping of numbers onto*  
8 *the actual registers of a given architecture. The mapping should be chosen to gain optimal*  
9 *density and should be shared by all users of a given architecture. It is recommended that*  
10 *this mapping be defined by the ABI authoring committee for each architecture.*

11 1. **DW\_OP\_reg0, DW\_OP\_reg1, ..., DW\_OP\_reg31**

12 The **DW\_OP\_reg<n>** operations encode the names of up to 32 registers,  
13 numbered from 0 through 31, inclusive. The object addressed is in register *n*.

14 2. **DW\_OP\_regx**

15 The **DW\_OP\_regx** operation has a single unsigned LEB128 literal operand  
16 that encodes the name of a register.

17 *These operations name a register location. To fetch the contents of a register, it is*  
18 *necessary to use one of the register based addressing operations, such as [DW\\_OP\\_bregx](#)*  
19 *(Section 2.5.1.2 on page 28).*

### 20 2.6.1.1.4 Implicit Location Descriptions

21 An implicit location description represents a piece or all of an object which has  
22 no actual location but whose contents are nonetheless either known or known to  
23 be undefined.

24 The following DWARF operations may be used to specify a value that has no  
25 location in the program but is a known constant or is computed from other  
26 locations and values in the program.

27 1. **DW\_OP\_implicit\_value**

28 The **DW\_OP\_implicit\_value** operation specifies an immediate value using  
29 two operands: an unsigned LEB128 length, followed by a sequence of bytes  
30 of the given length that contain the value.

31 2. **DW\_OP\_stack\_value**

32 The **DW\_OP\_stack\_value** operation specifies that the object does not exist in  
33 memory but its value is nonetheless known and is at the top of the DWARF  
34 expression stack. In this form of location description, the DWARF expression  
35 represents the actual value of the object, rather than its location. The  
36 **DW\_OP\_stack\_value** operation terminates the expression.

1     3. **DW\_OP\_implicit\_pointer**

2     *An optimizing compiler may eliminate a pointer, while still retaining the value that*  
3     *the pointer addressed. DW\_OP\_implicit\_pointer allows a producer to describe this*  
4     *value.*

5     The DW\_OP\_implicit\_pointer operation specifies that the object is a pointer  
6     that cannot be represented as a real pointer, even though the value it would  
7     point to can be described. In this form of location description, the DWARF  
8     expression refers to a debugging information entry that represents the actual  
9     value of the object to which the pointer would point. Thus, a consumer of the  
10    debug information would be able to show the value of the dereferenced  
11    pointer, even when it cannot show the value of the pointer itself.

12    The DW\_OP\_implicit\_pointer operation has two operands: a reference to a  
13    debugging information entry that describes the dereferenced object's value,  
14    and a signed number that is treated as a byte offset from the start of that  
15    value. The first operand is a 4-byte unsigned value in the 32-bit DWARF  
16    format, or an 8-byte unsigned value in the 64-bit DWARF format (see Section  
17    7.4 on page 204) that is used as the offset of a debugging information entry in  
18    the .debug\_info section of the current executable or shared object file. The  
19    second operand is a signed LEB128 number.

20    *The debugging information entry referenced by a DW\_OP\_implicit\_pointer*  
21    *operation is typically a DW\_TAG\_variable or DW\_TAG\_formal\_parameter entry*  
22    *whose DW\_AT\_location attribute gives a second DWARF expression or a location*  
23    *list that describes the value of the object, but the referenced entry may be any entry*  
24    *that contains a DW\_AT\_location or DW\_AT\_const\_value attribute (for example,*  
25    *DW\_TAG\_dwarf\_procedure). By using the second DWARF expression, a consumer*  
26    *can reconstruct the value of the object when asked to dereference the pointer described*  
27    *by the original DWARF expression containing the DW\_OP\_implicit\_pointer*  
28    *operation.*

29    DWARF location descriptions are intended to yield the **location** of a value rather than  
30    the value itself. An optimizing compiler may perform a number of code transformations  
31    where it becomes impossible to give a location for a value, but it remains possible to  
32    describe the value itself. Section 2.6.1.1.3 on page 39 describes operators that can be used  
33    to describe the location of a value when that value exists in a register but not in memory.  
34    The operations in this section are used to describe values that exist neither in memory nor  
35    in a single register.



1     **2.6.1.2 Composite Location Descriptions**

2     A composite location description describes an object or value which may be  
3     contained in part of a register or stored in more than one location. Each piece is  
4     described by a composition operation, which does not compute a value nor store  
5     any result on the DWARF stack. There may be one or more composition  
6     operations in a single composite location description. A series of such operations  
7     describes the parts of a value in memory address order.

8     Each composition operation is immediately preceded by a simple location  
9     description which describes the location where part of the resultant value is  
10    contained.

11    1. **DW\_OP\_piece**

12    The DW\_OP\_piece operation takes a single operand, which is an unsigned  
13    LEB128 number. The number describes the size in bytes of the piece of the  
14    object referenced by the preceding simple location description. If the piece is  
15    located in a register, but does not occupy the entire register, the placement of  
16    the piece within that register is defined by the ABI.

17    *Many compilers store a single variable in sets of registers, or store a variable partially*  
18    *in memory and partially in registers. DW\_OP\_piece provides a way of describing*  
19    *how large a part of a variable a particular DWARF location description refers to.*

20    2. **DW\_OP\_bit\_piece**

21    The DW\_OP\_bit\_piece operation takes two operands. The first is an unsigned  
22    LEB128 number that gives the size in bits of the piece. The second is an  
23    unsigned LEB128 number that gives the offset in bits from the location  
24    defined by the preceding DWARF location description.

25    Interpretation of the offset depends on the location description. If the location  
26    description is empty (see Section [2.6.1.1.1 on page 39](#)), the DW\_OP\_bit\_piece  
27    operation describes a piece consisting of the given number of bits whose  
28    values are undefined, and the offset is ignored. If the location is a memory  
29    address (see Section [2.6.1.1.2 on page 39](#)), the DW\_OP\_bit\_piece operation  
30    describes a sequence of bits relative to the location whose address is on the  
31    top of the DWARF stack using the bit numbering and direction conventions  
32    that are appropriate to the current language on the target system. In all other  
33    cases, the source of the piece is given by either a register location (see Section  
34    [2.6.1.1.3 on page 39](#)) or an implicit value description (see Section [2.6.1.1.4 on](#)  
35    [page 40](#)); the offset is from the least significant bit of the source value.

36    A composition operation that follows an empty location description indicates  
37    that the piece is undefined, for example because it has been optimized away.



1 *DW\_OP\_bit\_piece* is used instead of *DW\_OP\_piece* when the piece to be assembled into  
2 a value or assigned to is not byte-sized or is not at the start of a register or addressable  
3 unit of memory.

4 Whether or not a *DW\_OP\_piece* operation is equivalent to any *DW\_OP\_bit\_piece*  
5 operation with an offset of 0 is ABI dependent.

## 6 **2.6.2 Location Lists**

7 Location lists are used in place of location descriptions whenever the object  
8 whose location is being described can change location during its lifetime.  
9 Location lists are contained in a separate object file section called  
10 `.debug_loclists` or `.debug_loclists.dwo` (for split DWARF object files).

11 A location list is indicated by a location or other attribute whose value is of class  
12 `loclist` (see Section 7.5.5 on page 223).

13 *This location list representation, the `loclist` class, and the related `DW_AT_loclists_base`  
14 attribute are new in DWARF Version 5. Together they eliminate most or all of the object  
15 language relocations previously needed for location lists.*

16 A location list consists of a series of location list entries. Each location list entry is  
17 one of the following kinds:

- 18 • **Bounded location description.** This kind of entry provides a location  
19 description that specifies the location of an object that is valid over a  
20 lifetime bounded by a starting and ending address. The starting address is  
21 the lowest address of the address range over which the location is valid.  
22 The ending address is the address of the first location past the highest  
23 address of the address range. When the current PC is within the given  
24 range, the location description may be used to locate the specified object.

25 There are several kinds of bounded location description entries which  
26 differ in the way that they specify the starting and ending addresses.

27 The address ranges defined by the bounded location descriptions of a  
28 location list may overlap. When they do, they describe a situation in which  
29 an object exists simultaneously in more than one place. If all of the address  
30 ranges in a given location list do not collectively cover the entire range over  
31 which the object in question is defined, and there is no following default  
32 location description, it is assumed that the object is not available for the  
33 portion of the range that is not covered.

## Chapter 2. General Description

1 In the case of a bounded location description where the range is defined by  
2 a starting address and either an ending address or a length, a starting  
3 address consisting of the reserved address value (see Section 2.4.1 on  
4 [page 26](#)) indicates a non-existent range, which is equivalent to omitting the  
5 description.

- 6 • **Default location description.** This kind of entry provides a location  
7 description that specifies the location of an object that is valid when no  
8 bounded location description applies.
- 9 • **Base address.** This kind of entry provides an address to be used as the base  
10 address for beginning and ending address offsets given in certain kinds of  
11 bounded location description. The applicable base address of a bounded  
12 location description entry is the address specified by the closest preceding  
13 base address entry in the same location list. If there is no preceding base  
14 address entry, then the applicable base address defaults to the base address  
15 of the compilation unit (see Section 3.1.1 on [page 62](#)).

16 In the case of a compilation unit where all of the machine code is contained  
17 in a single contiguous section, no base address entry is needed.

18 If the base address is the reserved target address, either explicitly or by  
19 default, then the range of any bounded location description defined relative  
20 to that base address is non-existent, which is equivalent to omitting the  
21 description.

- 22 • **End-of-list.** This kind of entry marks the end of the location list.

23 A location list consists of a sequence of zero or more bounded location  
24 description or base address entries, optionally followed by a default location  
25 entry, and terminated by an end-of-list entry.

26 Each location list entry begins with a single byte identifying the kind of that  
27 entry, followed by zero or more operands depending on the kind.

28 In the descriptions that follow, these terms are used for operands:

- 29 • A **counted location description** operand consists of an unsigned ULEB  
30 integer giving the length of the location description (see Section 2.6.1 on  
31 [page 39](#)) that immediately follows.

## Chapter 2. General Description

- 1 • An **address index** operand is the index of an address in the `.debug_addr`  
2 section. This index is relative to the value of the `DW_AT_addr_base`  
3 attribute of the associated compilation unit. The address given by this kind  
4 of operand is not relative to the compilation unit base address.
- 5 • A **target address** operand is an address on the target machine. (Its size is  
6 the same as used for attribute values of class `address`, specifically,  
7 `DW_FORM_addr`.)

8 The following entry kinds are defined for use in both split or non-split units:

9 1. **DW\_LLE\_end\_of\_list**

10 An end-of-list entry contains no further data.

11 *A series of this kind of entry may be used for padding or alignment purposes.*

12 2. **DW\_LLE\_base\_addressx**

13 This is a form of base address entry that has one unsigned LEB128 operand.

14 The operand value is an address index (into the `.debug_addr` section) that  
15 indicates the applicable base address used by subsequent

16 `DW_LLE_offset_pair` entries.

17 3. **DW\_LLE\_startx\_endx**

18 This is a form of **bounded location description** entry (see page 43) that has  
19 two unsigned LEB128 operands. The operand values are address indices (into  
20 the `.debug_addr` section). These indicate the starting and ending addresses,  
21 respectively, that define the address range for which this location is valid.

22 These operands are followed by a counted location description.

23 4. **DW\_LLE\_startx\_length**

24 This is a form of **bounded location description** entry (see page 43) that has  
25 two unsigned LEB128 operands. The first value is an address index (into the  
26 `.debug_addr` section) that indicates the beginning of the address range over  
27 which the location is valid. The second value is the length of the range. These  
28 operands are followed by a counted location description.

29 5. **DW\_LLE\_offset\_pair**

30 This is a form of **bounded location description** entry (see page 43) that has  
31 two unsigned LEB128 operands. The values of these operands are the starting  
32 and ending offsets, respectively, relative to the applicable base address, that  
33 define the address range for which this location is valid. These operands are  
34 followed by a counted location description.

## Chapter 2. General Description

1     6. **DW\_LLE\_default\_location**

2     The operand is a counted location description which defines where an object  
3     is located if no prior location description is valid.

4     7. **DW\_LLE\_include\_loclistx**

5     This is a form of list inclusion, that has one unsigned LEB128 operand. The  
6     value is an index into the `.debug_loclists` section, interpreted the same way  
7     as the operand of [DW\\_FORM\\_loclistx](#) to find a target list of entries, which  
8     will be regarded as part of the current location list, up to the  
9     [DW\\_LLE\\_end\\_of\\_list](#) entry.

10    The following kinds of location list entries are defined for use only in non-split  
11    DWARF units:

12    7. **DW\_LLE\_base\_address**

13    A base address entry has one target address operand. This address is used as  
14    the base address when interpreting offsets in subsequent location list entries  
15    of kind [DW\\_LLE\\_offset\\_pair](#).

16    8. **DW\_LLE\_start\_end**

17    This is a form of [bounded location description](#) entry (see page 43) that has  
18    two target address operands. These indicate the starting and ending  
19    addresses, respectively, that define the address range for which the location is  
20    valid. These operands are followed by a counted location description.

21    9. **DW\_LLE\_start\_length**

22    This is a form of [bounded location description](#) entry (see page 43) that has  
23    one target address operand value and an unsigned LEB128 integer operand  
24    value. The address is the beginning address of the range over which the  
25    location description is valid, and the length is the number of bytes in that  
26    range. These operands are followed by a counted location description.

27    10. **DW\_LLE\_include\_loclist**

28    This is a form of list inclusion, that has one offset operand. The value is an  
29    offset into the `.debug_loclists` section, like the operand of  
30    [DW\\_FORM\\_sec\\_offset](#). The offset identifies the first entry of a location list  
31    whose entries are to be regarded as part of the current location list, up to the  
32    [DW\\_LLE\\_end\\_of\\_list](#) entry.

## 2.7 Types of Program Entities

Any debugging information entry describing a declaration that has a type has a `DW_AT_type` attribute, whose value is a reference to another debugging information entry. The entry referenced may describe a base type, that is, a type that is not defined in terms of other data types, or it may describe a user-defined type, such as an array, structure or enumeration. Alternatively, the entry referenced may describe a type modifier, such as constant, packed, pointer, reference or volatile, which in turn will reference another entry describing a type or type modifier (using a `DW_AT_type` attribute of its own). See Chapter 5 following for descriptions of the entries describing base types, user-defined types and type modifiers.

## 2.8 Accessibility of Declarations

*Some languages, notably C++ and Ada, have the concept of the accessibility of an object or of some other program entity. The accessibility specifies which classes of other program objects are permitted access to the object in question.*

The accessibility of a declaration is represented by a `DW_AT_accessibility` attribute, whose value is a constant drawn from the set of codes listed in Table 2.4.

Table 2.4: Accessibility codes

---

`DW_ACCESS_public`  
`DW_ACCESS_private`  
`DW_ACCESS_protected`

---

## 2.9 Visibility of Declarations

*Several languages (such as Modula-2) have the concept of the visibility of a declaration. The visibility specifies which declarations are to be visible outside of the entity in which they are declared.*

The visibility of a declaration is represented by a `DW_AT_visibility` attribute, whose value is a constant drawn from the set of codes listed in Table 2.5 on the next page.

Table 2.5: Visibility codes

---

DW\_VIS\_local  
DW\_VIS\_exported  
DW\_VIS\_qualified

---

## 2.10 Virtuality of Declarations

C++ provides for virtual and pure virtual structure or class member functions and for virtual base classes.

The virtuality of a declaration is represented by a `DW_AT_virtuality` attribute, whose value is a constant drawn from the set of codes listed in Table 2.6.

Table 2.6: Virtuality codes

---

DW\_VIRTUALITY\_none  
DW\_VIRTUALITY\_virtual  
DW\_VIRTUALITY\_pure\_virtual

---

## 2.11 Artificial Entries

A compiler may wish to generate debugging information entries for objects or types that were not actually declared in the source of the application. An example is a formal parameter entry to represent the hidden `this` parameter that most C++ implementations pass as the first argument to non-static member functions.

Any debugging information entry representing the declaration of an object or type artificially generated by a compiler and not explicitly declared by the source program may have a `DW_AT_artificial` attribute, which is a flag.

## 2.12 Address Classes

Some systems support different classes of addresses. The address class may affect the way a pointer is dereferenced or the way a subroutine is called.

Any debugging information entry representing a pointer or reference type or a subroutine or subroutine type may have a `DW_AT_address_class` attribute, whose value is an integer constant. The set of permissible values is specific to each target architecture. The value `DW_ADDR_none`, however, is common to all encodings, and means that no address class has been specified.

## 2.13 Non-Defining Declarations and Completions

A debugging information entry representing a program entity typically represents the defining declaration of that entity. In certain contexts, however, a debugger might need information about a declaration of an entity that is not also a definition, or is otherwise incomplete, to evaluate an expression correctly.

As an example, consider the following fragment of C code:

```
void myfunc()
{
    int x;
    {
        extern float x;
        g(x);
    }
}
```

C scoping rules require that the value of the variable `x` passed to the function `g` is the value of the global `float` variable `x` rather than of the local `int` variable `x`.

### 2.13.1 Non-Defining Declarations

A debugging information entry that represents a non-defining or otherwise incomplete declaration of a program entity has a `DW_AT_declaration` attribute, which is a [flag](#).

A non-defining type declaration may nonetheless have children as illustrated in [Section E.2.3 on page 408](#).

### 2.13.2 Declarations Completing Non-Defining Declarations

A debugging information entry that represents a declaration that completes another (earlier) non-defining declaration may have a `DW_AT_specification` attribute whose value is a [reference](#) to the debugging information entry representing the non-defining declaration. A debugging information entry with a `DW_AT_specification` attribute does not need to duplicate information provided by the debugging information entry referenced by that specification attribute.

When the non-defining declaration is contained within a type that has been placed in a separate type unit (see [Section 3.1.4 on page 71](#)), the `DW_AT_specification` attribute cannot refer directly to the entry in the type unit. Instead, the current compilation unit may contain a “skeleton” declaration of the type, which contains only the relevant declaration and its ancestors as necessary to provide the context (including containing types and namespaces). The `DW_AT_specification` attribute would then be a reference to the declaration entry within the skeleton declaration tree. The debugging information entry for the top-level type in the skeleton tree may contain a `DW_AT_signature` attribute whose value is the type signature (see [Section 7.32 on page 257](#)).

Not all attributes of the debugging information entry referenced by a `DW_AT_specification` attribute apply to the referring debugging information entry. For example, `DW_AT_sibling` and `DW_AT_declaration` cannot apply to a referring entry.

## 2.14 Declaration Coordinates

*It is sometimes useful in a debugger to be able to associate a declaration with its occurrence in the program source.*

Any debugging information entry representing the declaration of an object, module, subprogram or type may have `DW_AT_decl_file`, `DW_AT_decl_line` and `DW_AT_decl_column` attributes, each of whose value is an unsigned [integer constant](#).

The value of the `DW_AT_decl_file` attribute corresponds to a file number from the line number information table for the compilation unit containing the debugging information entry and represents the source file in which the declaration appeared (see [Section 6.2 on page 155](#)).

The value of the `DW_AT_decl_line` attribute represents the source line number at which the first character of the identifier of the declared object appears. The value 0 indicates that no source line has been specified.



1 The value of the `DW_AT_decl_column` attribute represents the source column  
2 number at which the first character of the identifier of the declared object  
3 appears. The value 0 indicates that no column has been specified.

## 4 **2.15 Identifier Names**

5 Any debugging information entry representing a program entity that has been  
6 given a name may have a `DW_AT_name` attribute, whose value of class `string`  
7 represents the name. A debugging information entry containing no name  
8 attribute, or containing a name attribute whose value consists of a name  
9 containing a single null byte, represents a program entity for which no name was  
10 given in the source.

11 *Because the names of program objects described by DWARF are the names as they appear*  
12 *in the source program, implementations of language translators that use some form of*  
13 *mangled name (as do many implementations of C++) should use the unmangled form of*  
14 *the name in the `DW_AT_name` attribute, including the keyword operator (in names such*  
15 *as “operator +”), if present. See also Section 2.22 following regarding the use of*  
16 *`DW_AT_linkage_name` for mangled names. Sequences of multiple whitespace characters*  
17 *may be compressed.*

18 *For additional discussion, see the Best Practices section of the DWARF Wiki*  
19 *([http://wiki.dwarfstd.org/index.php?title=Best\\_Practices](http://wiki.dwarfstd.org/index.php?title=Best_Practices).)*

## 20 **2.16 Data Locations and DWARF Procedures**

21 Any debugging information entry describing a data object (which includes  
22 variables and parameters) or `common blocks` may have a `DW_AT_location`  
23 attribute, whose value is a location description (see Section 2.6 on page 38).

24 A DWARF procedure is represented by any debugging information entry that  
25 has a `DW_AT_location` attribute. If a suitable entry is not otherwise available, a  
26 DWARF procedure can be represented using a debugging information entry with  
27 the tag `DW_TAG_dwarf_procedure` together with a `DW_AT_location` attribute.

28 A DWARF procedure is called by a `DW_OP_call2`, `DW_OP_call4` or  
29 `DW_OP_call_ref` DWARF expression operator (see Section 2.5.1.5 on page 35).

## 2.17 Code Addresses, Ranges and Base Addresses

Any debugging information entry describing an entity that has a machine code address or range of machine code addresses, which includes compilation units, module initialization, subroutines, lexical blocks, try/catch blocks (see Section 3.8 on page 97), labels and the like, may have

- A `DW_AT_low_pc` attribute for a single address,
- A `DW_AT_low_pc` and `DW_AT_high_pc` pair of attributes for a single contiguous range of addresses, or
- A `DW_AT_ranges` attribute for a non-contiguous range of addresses.

If a producer emits no machine code for an entity, none of these attributes are specified. Equivalently, a producer may emit such an attribute using the reserved target address (see Section 2.4.1 on page 26) for the non-existent entity.

The **base address** of the scope for any of the debugging information entries listed above is given by either the `DW_AT_low_pc` attribute or the first address in the first range entry in the list of ranges given by the `DW_AT_ranges` attribute. If there is no such attribute, the base address is undefined.

### 2.17.1 Single Address

When there is a single address associated with an entity, such as a label or alternate entry point of a subprogram, the entry has a `DW_AT_low_pc` attribute whose value is the address for the entity.

## 2.17.2 Contiguous Address Range

When the set of addresses of a debugging information entry can be described as a single contiguous range, the entry may have a `DW_AT_low_pc` and `DW_AT_high_pc` pair of attributes. The value of the `DW_AT_low_pc` attribute is the address of the first instruction associated with the entity. If the value of the `DW_AT_high_pc` is of class address, it is the address of the first location past the last instruction associated with the entity; if it is of class constant, the value is an unsigned integer offset which when added to the low PC gives the address of the first location past the last instruction associated with the entity.

*The high PC value may be beyond the last valid instruction in the executable.*

## 2.17.3 Non-Contiguous Address Ranges

Range lists are used when the set of addresses for a debugging information entry cannot be described as a single contiguous range. Range lists are contained in a separate object file section called `.debug_rnglists` or `.debug_rnglists.dwo` (in split units).

A range list is identified by a `DW_AT_ranges` or other attribute whose value is of class `rnglist` (see Section 7.5.5 on page 223).

*This range list representation, the `rnglist` class, and the related `DW_AT_rnglists_base` attribute are new in DWARF Version 5. Together they eliminate most or all of the object language relocations previously needed for range lists.*

Each range list entry is one of the following kinds:

- **Bounded range.** This kind of entry defines an address range that is included in the range list. The starting address is the lowest address of the address range. The ending address is the address of the first location past the highest address of the address range.

There are several kinds of bounded range entries which specify the starting and ending addresses in different ways.

In the case of a range list entry where the range is defined by a starting address and either an ending address or a length, a starting address consisting of the reserved address value (see Section 2.4.1 on page 26) indicates a non-existent range, which is equivalent to omitting the description.

## Chapter 2. General Description

- **Base address.** This kind of entry provides an address to be used as the base address for the beginning and ending address offsets given in certain bounded range entries. The applicable base address of a range list entry is determined by the closest preceding base address entry in the same range list. If there is no preceding base address entry, then the applicable base address defaults to the base address of the compilation unit (see Section 3.1.1 on page 62).

In the case of a compilation unit where all of the machine code is contained in a single contiguous section, no base address entry is needed.

If the base address is the reserved target address, either explicitly or by default, then the range of any range list entry defined relative to that base address is non-existent, which is equivalent to omitting the range list entry.

- **End-of-list.** This kind of entry marks the end of the range list.

Each range list consists of a sequence of zero or more bounded range or base address entries, terminated by an end-of-list entry.

A range list containing only an end-of-list entry describes an empty scope (which contains no instructions).

Bounded range entries in a range list may not overlap. There is no requirement that the entries be ordered in any particular way.

*A bounded range entry whose beginning and ending addresses are equal (including zero) indicates an empty range and may be ignored.*

Each range list entry begins with a single byte identifying the kind of that entry, followed by zero or more operands depending on the kind.

In the descriptions that follow, the term **address index** means the index of an address in the `.debug_addr` section. This index is relative to the value of the `DW_AT_addr_base` attribute of the associated compilation unit. The address given by this kind of operand is *not* relative to the compilation unit base address.

The following entry kinds are defined for use in both split or non-split units:

1. **DW\_RLE\_end\_of\_list**

An end-of-list entry contains no further data.

*A series of this kind of entry may be used for padding or alignment purposes.*

2. **DW\_RLE\_base\_addressx**

A base address entry has one unsigned LEB128 operand. The operand value is an address index (into the `.debug_addr` section) that indicates the applicable base address used by following `DW_RLE_offset_pair` entries.

## Chapter 2. General Description

### 3. **DW\_RLE\_startx\_endx**

This is a form of [bounded range](#) (see page 53) entry that has two unsigned LEB128 operands. The operand values are address indices (into the `.debug_addr` section) that indicate the starting and ending addresses, respectively, that define the address range.

### 4. **DW\_RLE\_startx\_length**

This is a form of [bounded range](#) (see page 53) entry that has two unsigned ULEB operands. The first value is an address index (into the `.debug_addr` section) that indicates the beginning of the address range. The second value is the length of the range.

### 5. **DW\_RLE\_offset\_pair**

This is a form of [bounded range](#) (see page 53) entry that has two unsigned LEB128 operands. The values of these operands are the starting and ending offsets, respectively, relative to the applicable base address, that define the address range.

### 6. **DW\_RLE\_include\_rnglistx**

This is a form of range inclusion, that has one unsigned LEB128 operand. The value is an index into the `.debug_rnglists` section, interpreted the same way as the operand of [DW\\_FORM\\_rnglistx](#) to find a target list of entries, which will be regarded as part of the current range list, up to the [DW\\_RLE\\_end\\_of\\_list](#) entry.

The following kinds of range entry may be used only in non-split units:

### 6. **DW\_RLE\_base\_address**

A base address entry has one target address operand. This operand is the same size as used in [DW\\_FORM\\_addr](#). This address is used as the base address when interpreting offsets in subsequent location list entries of kind [DW\\_RLE\\_offset\\_pair](#).

### 7. **DW\_RLE\_start\_end**

This is a form of [bounded range](#) (see page 53) entry that has two target address operands. Each operand is the same size as used in [DW\\_FORM\\_addr](#). These indicate the starting and ending addresses, respectively, that define the address range for which the following location is valid.

1 8. **DW\_RLE\_start\_length**

2 This is a form of [bounded range](#) (see page 53) entry that has one target  
3 address operand value and an unsigned LEB128 integer length operand  
4 value. The address is the beginning address of the range over which the  
5 location description is valid, and the length is the number of bytes in that  
6 range.

7 9. **DW\_RLE\_include\_rnglist**

8 This is a form of list inclusion, that has one offset operand. The value is an  
9 offset into the `.debug_rnglists` section, like the operand of a  
10 [DW\\_FORM\\_sec\\_offset](#) location list. The offset identifies the first entry of a  
11 location list whose entries are to be regarded as part of the current location  
12 list, up to the [DW\\_RLE\\_end\\_of\\_list](#) entry.

13 **2.18 Entry Address**

14 *The entry or first executable instruction generated for an entity, if applicable, is often the*  
15 *lowest addressed instruction of a contiguous range of instructions. In other cases, the*  
16 *entry address needs to be specified explicitly.*

17 Any debugging information entry describing an entity that has a range of code  
18 addresses, which includes compilation units, module initialization, subroutines,  
19 [lexical blocks](#), [try/catch blocks](#), and the like, may have a [DW\\_AT\\_entry\\_pc](#)  
20 attribute to indicate the **entry address** which is the address of the instruction  
21 where execution begins within that range of addresses. If the value of the  
22 `DW_AT_entry_pc` attribute is of class [address](#) that address is the entry address;  
23 or, if it is of class [constant](#), the value is an unsigned integer offset which, when  
24 added to the base address of the function, gives the entry address.

25 If no `DW_AT_entry_pc` attribute is present, then the entry address is assumed to  
26 be the same as the base address of the containing scope.

27 **2.19 Static and Dynamic Values of Attributes**

28 Some attributes that apply to types specify a property (such as the lower bound  
29 of an array) that is an integer value, where the value may be known during  
30 compilation or may be computed dynamically during execution.

## Chapter 2. General Description

1 The value of these attributes is determined based on the class as follows:

- 2 • For a **constant**, the value of the constant is the value of the attribute.
- 3 • For a **reference**, the value of the attribute is determined indirectly via a  
4 reference to another debugging information entry.
  - 5 – If the referenced entry describes a constant (e.g., has a  
6 **DW\_AT\_const\_value** attribute), the attribute value is the value of that  
7 constant.
  - 8 – If the referenced entry describes a data object (see Section 4.1 on  
9 page 99) or common block (see Section 4.2 on page 102), the attribute  
10 value is the value of the data object or common block.
  - 11 – If the referenced entry represents a data member (e.g. has either a  
12 **DW\_AT\_data\_member\_location** or a **DW\_AT\_data\_bit\_offset**  
13 attribute), the attribute value is the value of the data member.
- 14 • For an **exprval**, the value is interpreted as a DWARF expression; evaluation  
15 of the expression yields the value of the attribute.

16 *Prior to DWARF Version 6, a reference to a DWARF procedure (see Section 2.16 on*  
17 *page 51) that is not a data object or common block was allowed. This type of reference*  
18 *was removed in DWARF Version 6. Instead, a producer may use a form of class **exprval***  
19 *or **locdesc** with a **DW\_OP\_call\_ref** operator to call the DWARF procedure.*

### 20 2.20 Entity Descriptions

21 *Some debugging information entries may describe entities in the program that are*  
22 *artificial, or which otherwise have a “name” that is not a valid identifier in the*  
23 *programming language. This attribute provides a means for the producer to indicate the*  
24 *purpose or usage of the containing debugging infor*

25 Generally, any debugging information entry that has, or may have, a  
26 **DW\_AT\_name** attribute, may also have a **DW\_AT\_description** attribute whose  
27 value is a null-terminated string providing a description of the entity.

28 *It is expected that a debugger will display these descriptions as part of displaying other*  
29 *properties of an entity.*



## 2.21 Byte and Bit Sizes

Many debugging information entries allow either a `DW_AT_byte_size` attribute or a `DW_AT_bit_size` attribute, whose [integer constant](#) value (see Section 2.19) specifies an amount of storage. The value of the `DW_AT_byte_size` attribute is interpreted in bytes and the value of the `DW_AT_bit_size` attribute is interpreted in bits. The `DW_AT_string_length_byte_size` and `DW_AT_string_length_bit_size` attributes are similar.

In addition, the [integer constant](#) value of a `DW_AT_byte_stride` attribute is interpreted in bytes and the [integer constant](#) value of a `DW_AT_bit_stride` attribute is interpreted in bits.

## 2.22 Linkage Names

*Some language implementations, notably C++ and similar languages, make use of implementation-defined names within object files that are different from the identifier names (see Section 2.15 on page 51) of entities as they appear in the source. Such names, sometimes known as mangled names, are used in various ways, such as: to encode additional information about an entity, to distinguish multiple entities that have the same name, and so on. When an entity has an associated distinct linkage name it may sometimes be useful for a producer to include this name in the DWARF description of the program to facilitate consumer access to and use of object file information about an entity and/or information that is encoded in the linkage name itself.*

A debugging information entry may have a `DW_AT_linkage_name` attribute whose value is a null-terminated string containing the object file linkage name associated with the corresponding entity.

## 2.23 Template Parameters

*In C++, a template is a generic definition of a class, function, member function, or typedef (alias). A template has formal parameters that can be types or constant values; the class, function, member function, or typedef is instantiated differently for each distinct combination of type or value actual parameters. DWARF does not represent the generic template definition, but does represent each instantiation.*

A debugging information entry that represents a template instantiation will contain child entries describing the actual template parameters. The containing entry and each of its child entries reference a template parameter entry in any



1 circumstance where the template definition referenced a formal template  
2 parameter.

3 A template type parameter is represented by a debugging information entry with  
4 the tag `DW_TAG_template_type_parameter`. A template value parameter is  
5 represented by a debugging information entry with the tag  
6 `DW_TAG_template_value_parameter`. The actual template parameter entries  
7 appear in the same order as the corresponding template formal parameter  
8 declarations in the source program.

9 A type or value parameter entry may have a `DW_AT_name` attribute, whose  
10 value is a null-terminated string containing the name of the corresponding  
11 formal parameter. The entry may also have a `DW_AT_default_value` attribute,  
12 which is a flag indicating that the value corresponds to the default argument for  
13 the template parameter.

14 A template type parameter entry has a `DW_AT_type` attribute describing the  
15 actual type by which the formal is replaced.

16 A template value parameter entry has a `DW_AT_type` attribute describing the  
17 type of the parameterized value. The entry also has an attribute giving the actual  
18 compile-time or run-time constant value of the value parameter for this  
19 instantiation. This can be a `DW_AT_const_value` attribute, whose value is the  
20 compile-time constant value as represented on the target architecture, or a  
21 `DW_AT_location` attribute, whose value is a single location description for the  
22 run-time constant address.

### 23 **2.24 Alignment**

24 A debugging information entry may have a `DW_AT_alignment` attribute whose  
25 value of class `constant` is a positive, non-zero, integer describing the alignment of  
26 the entity.

27 *For example, an alignment attribute whose value is 8 indicates that the entity to which it*  
28 *applies occurs at an address that is a multiple of eight (not a multiple of 2<sup>8</sup> or 256).*

## Chapter 2. General Description

*(empty page)*

# Chapter 3

## Program Scope Entries

This section describes debugging information entries that relate to different levels of program scope: compilation, module, subprogram, and so on. Except for separate type entries (see Section 3.1.4 on page 71), these entries may be thought of as ranges of text addresses within the program.

### 3.1 Unit Entries

A DWARF object file is an object file that contains one or more DWARF compilation units, of which there are these kinds:

- A **full compilation unit** describes a complete compilation, possibly in combination with related partial compilation units and/or type units.
- A **partial compilation unit** describes a part of a compilation (generally corresponding to an imported module) which is imported into one or more related full compilation units.
- A **type unit** is a specialized unit (similar to a compilation unit) that represents a type whose description may be usefully shared by multiple other units.

*These first three kinds of compilation unit are sometimes called “conventional” compilation units—they are kinds of compilation units that were defined prior to DWARF Version 5. Conventional compilation units are part of the same object file as the compiled code and data (whether relocatable, executable, shared and so on). The word “conventional” is usually omitted in these names, unless needed to distinguish them from the similar split compilation units below.*

## Chapter 3. Program Scope Entries

- A **skeleton compilation unit** represents the DWARF debugging information for a compilation using a minimal description that identifies a separate split compilation unit that provides the remainder (and most) of the description.

*A skeleton compilation acts as a minimal conventional full compilation (see above) that identifies and is paired with a corresponding split full compilation (as described below). Like the conventional compilation units, a skeleton compilation unit is part of the same object file as the compiled code and data.*

- A **split compilation unit** describes a complete compilation, possibly in combination with related type compilation units. It corresponds to a specific skeleton compilation unit.
- A **split type unit** is a specialized compilation unit that represents a type whose description may be usefully shared by multiple other units.

*Split compilation units and split type units may be contained in object files separate from those containing the program code and data. These object files are not processed by a linker; thus, split units do not depend on underlying object file relocations.*

*Either a full compilation unit or a partial compilation unit may be logically incorporated into another compilation unit using an imported unit entry (see Section 3.2.5 on page 78).*

*A partial compilation unit is not defined for use within a split object file.*

*In the remainder of this document, the word “compilation” in the phrase “compilation unit” is generally omitted, unless it is deemed needed for clarity or emphasis.*

### 3.1.1 Full and Partial Compilation Unit Entries

A full compilation unit is represented by a debugging information entry with the tag **DW\_TAG\_compile\_unit**. A partial compilation unit is represented by a debugging information entry with the tag **DW\_TAG\_partial\_unit**.

In a simple compilation, a single compilation unit with the tag **DW\_TAG\_compile\_unit** represents a complete object file and the tag **DW\_TAG\_partial\_unit** (as well as tag **DW\_TAG\_type\_unit**) is not used. In a compilation employing the DWARF space compression and duplicate elimination techniques from Appendix E.1 on page 386, multiple compilation units using the tags **DW\_TAG\_compile\_unit**, **DW\_TAG\_partial\_unit** and/or **DW\_TAG\_type\_unit** are used to represent portions of an object file.

## Chapter 3. Program Scope Entries

1 *A full compilation unit typically represents the text and data contributed to an*  
2 *executable by a single relocatable object file. It may be derived from several source files,*  
3 *including pre-processed header files. A partial compilation unit typically represents a*  
4 *part of the text and data of a relocatable object file, in a manner that can potentially be*  
5 *shared with the results of other compilations to save space. It may be derived from an*  
6 *“include file,” template instantiation, or other implementation-dependent portion of a*  
7 *compilation. A full compilation unit can also function in a manner similar to a partial*  
8 *compilation unit in some cases. See Appendix E on page 386 for discussion of related*  
9 *compression techniques.*

10 A full or partial compilation unit entry owns debugging information entries that  
11 represent all or part of the declarations made in the corresponding compilation.  
12 In the case of a partial compilation unit, the containing scope of its owned  
13 declarations is indicated by imported unit entries in one or more other  
14 compilation unit entries that refer to that partial compilation unit (see Section  
15 3.2.5 on page 78).

16 A full or partial compilation unit entry may have the following attributes:

- 17 1. Either a `DW_AT_low_pc` and `DW_AT_high_pc` pair of attributes or a  
18 `DW_AT_ranges` attribute whose values encode the contiguous or  
19 non-contiguous address ranges, respectively, of the machine instructions  
20 generated for the compilation unit (see Section 2.17 on page 52).

21 A `DW_AT_low_pc` attribute may also be specified in combination with  
22 `DW_AT_ranges` to specify the default base address for use in location lists  
23 (see Section 2.6.2 on page 43) and range lists (see Section 2.17.3 on page 53).

- 24 2. A `DW_AT_name` attribute whose value is a null-terminated string containing  
25 the full or relative path name (relative to the value of the `DW_AT_comp_dir`  
26 attribute, see below) of the primary source file from which the compilation  
27 unit was derived.

- 28 3. A `DW_AT_language_name` attribute whose constant value is an integer code  
29 indicating the source language of the compilation unit. The set of language  
30 names and their meanings are given in Table 3.1 on the following page.

31 *The most recent list of approved language names and applicable versions may be*  
32 *found at <http://dwarfstd.org/LanguagesV6.php>.*

## Chapter 3. Program Scope Entries

Table 3.1: Language names

Language name	Meaning	Version Scheme (See Table 3.2)
DW_LNAME_Ada †	ISO Ada	YYYY
DW_LNAME_Assembly	Assembly	
DW_LNAME_BLISS	BLISS	
DW_LNAME_C	ISO C	YYYYMM
DW_LNAME_C_plus_plus	ISO C++	YYYYMM
DW_LNAME_Cobol	ISO COBOL	YYYY
DW_LNAME_CPP_for_OpenCL †	C++ for OpenCL	VVMM
DW_LNAME_Crystal	Crystal	
DW_LNAME_C_sharp	C#	
DW_LNAME_D †	D	
DW_LNAME_Dylan †	Dylan	
DW_LNAME_Fortran	ISO Fortran	YYYY
DW_LNAME_Go †	Go	
DW_LNAME_GLSL †	OpenGL Shading Language	VVMMPP
DW_LNAME_GLSL_ES †	OpenGL ES Shading Language	VVMMPP
DW_LNAME_Haskell †	Haskell	
DW_LNAME_HLSL †	High-Level Shading Language	YYYY
DW_LNAME_Java	Java	
DW_LNAME_Julia †	Julia	
DW_LNAME_Kotlin	Kotlin	
DW_LNAME_Modula2	ISO Modula-2	
DW_LNAME_Modula3	Modula-3	
DW_LNAME_Mojo	Mojo	
DW_LNAME_ObjC	Objective C	YYYYMM
DW_LNAME_ObjC_plus_plus	Objective C++	YYYYMM
DW_LNAME_OCaml †	OCaml	

*Continued on next page*

## Chapter 3. Program Scope Entries

Language name	Meaning	Version Scheme
<code>DW_LNAME_OpenCL_C</code> <sup>1</sup> †	OpenCL C	VVMM
<code>DW_LNAME_OpenCL_CPP</code> †	OpenCL C++	VVMM
<code>DW_LNAME_Pascal</code>	ISO Pascal	YYYY
<code>DW_LNAME_PLI</code> †	ANSI PL/I	
<code>DW_LNAME_Python</code> †	Python	
<code>DW_LNAME_RenderScript</code> †	RenderScript Kernel Language	
<code>DW_LNAME_Ruby</code> †	Ruby	VVMM
<code>DW_LNAME_Rust</code> †	Rust	
<code>DW_LNAME_Swift</code>	Swift	
<code>DW_LNAME_SYCL</code> †	SYCL	YYYYRR
<code>DW_LNAME_UPC</code>	UPC (Unified Parallel C)	
<code>DW_LNAME_Zig</code>	Zig	

† Support for these languages is limited

- 1 4. A `DW_AT_language_version` attribute may be specified whose constant  
2 value is an integer value that indicates the version of the source language.  
3 This value is encoded using one of several schemes as shown in Table 3.2 on  
4 the next page. A value of zero is equivalent to omitting this attribute.
- 5 5. A `DW_AT_stmt_list` attribute whose value is a section offset to the line  
6 number information for this compilation unit.
- 7 This information is placed in a separate object file section from the debugging  
8 information entries themselves. The value of the statement list attribute is the  
9 offset in the `.debug_line` section of the first byte of the line number  
10 information for this compilation unit (see Section 6.2 on page 155).
- 11 6. A `DW_AT_macros` attribute whose value is a section offset to the macro  
12 information for this compilation unit.
- 13 This information is placed in a separate object file section from the debugging  
14 information entries themselves. The value of the macro information attribute  
15 is the offset in the `.debug_macro` section of the first byte of the macro  
16 information for this compilation unit (see Section 6.3 on page 172).
- 17 *The `DW_AT_macros` attribute is new in DWARF Version 5, and supersedes the*  
18 *`DW_AT_macro_info` attribute of earlier DWARF versions. While `DW_AT_macros`*

<sup>1</sup>This is equivalent to `DW_LANG_OpenCL` in DWARF Version 5

## Chapter 3. Program Scope Entries

Table 3.2: Version Encoding Schemes

Scheme	Encoding
YYYY	Year in which the language definition was released
YYYYMM	Year in which the language definition was released times 100 plus the ordinal number of the month (from 1 to 12). <i>For example, 202206 represents June of 2022.</i>
YYYYRR	Year in which the language definition was released times 100 plus the revision number. <i>For example, 202007 represents version 2020 revision 7 while 202011 represents version 2020 revision 11.</i>
VVMM	Major version number times 100 plus the minor version number. <i>For example, 306 represents version 3.6 while 312 represents version 3.12.</i>
VVMMPP	Major version number times 10,000 plus the minor version number times 100 plus the patch version number. <i>For example, 30607 represents version 3.6.7 while 31215 represents version 3.12.15.</i>

1 *and DW\_AT\_macro\_info attributes cannot both occur in the same compilation unit,*  
2 *both may be found in the set of units that make up an executable or shared object file.*  
3 *The two attributes have distinct encodings to facilitate such coexistence.*

- 4 7. A **DW\_AT\_comp\_dir** attribute whose value is a null-terminated string  
5 containing the current working directory of the compilation command that  
6 produced this compilation unit in whatever form makes sense for the host  
7 system.

8 If a relative path is used in DW\_AT\_comp\_dir, it will be relative to the  
9 location of the linked image containing the DW\_AT\_comp\_dir entry.

10 *In some cases a producer may allow the user to specify a relative path for*  
11 *[DW\\_AT\\_comp\\_dir](#). There are a few cases in which this is useful, but in general using*  
12 *a relative path for [DW\\_AT\\_comp\\_dir](#) is discouraged as it will not work well in many*  
13 *cases including the following: different relative paths are used within the same build;*  
14 *the build system creates multiple linked images in different directories; the final linked*  
15 *image is moved before being debugged; .o files that need to be debugged directly.*



## Chapter 3. Program Scope Entries

- 1 8. A **DW\_AT\_producer** attribute whose value is a null-terminated string  
2 containing information about the compiler that produced the compilation  
3 unit.
- 4 *The actual contents of the string will be specific to each producer, but should begin*  
5 *with the name of the compiler vendor or some other identifying character sequence*  
6 *that will avoid confusion with other producer values.*
- 7 9. A **DW\_AT\_identifier\_case** attribute whose integer constant value is a code  
8 describing the treatment of identifiers within this compilation unit. The set of  
9 identifier case codes is given in Table 3.3.

Table 3.3: Identifier case codes

---

DW\_ID\_case\_sensitive  
DW\_ID\_up\_case  
DW\_ID\_down\_case  
DW\_ID\_case\_insensitive

---

10 **DW\_ID\_case\_sensitive** is the default for all compilation units that do not  
11 have this attribute. It indicates that names given as the values of  
12 **DW\_AT\_name** attributes in debugging information entries for the  
13 compilation unit reflect the names as they appear in the source program.

14 *A debugger should be sensitive to the case of identifier names when doing identifier*  
15 *lookups.*

16 **DW\_ID\_up\_case** means that the producer of the debugging information for  
17 this compilation unit converted all source names to upper case. The values of  
18 the name attributes may not reflect the names as they appear in the source  
19 program.

20 *A debugger should convert all names to upper case when doing lookups.*

21 **DW\_ID\_down\_case** means that the producer of the debugging information  
22 for this compilation unit converted all source names to lower case. The values  
23 of the name attributes may not reflect the names as they appear in the source  
24 program.

25 *A debugger should convert all names to lower case when doing lookups.*

## Chapter 3. Program Scope Entries

1 **DW\_ID\_case\_insensitive** means that the values of the name attributes reflect  
2 the names as they appear in the source program but that case is not  
3 significant.

4 *A debugger should ignore case when doing lookups.*

- 5 10. A **DW\_AT\_base\_types** attribute whose value is a **reference**. This attribute  
6 points to a debugging information entry representing another compilation  
7 unit. It may be used to specify the compilation unit containing the base type  
8 entries used by entries in the current compilation unit (see Section 5.1 on  
9 [page 106](#)).

10 *This attribute provides a consumer a way to find the definition of base types for a*  
11 *compilation unit that does not itself contain such definitions. This allows a consumer,*  
12 *for example, to interpret a type conversion to a base type correctly.*

- 13 11. A **DW\_AT\_use\_UTF8** attribute, which is a **flag** whose presence indicates that  
14 all strings (such as the names of declared entities in the source program, or  
15 filenames in the line number table) are represented using the UTF-8  
16 representation.

- 17 12. A **DW\_AT\_main\_subprogram** attribute, which is a **flag**, whose presence  
18 indicates that the compilation unit contains a subprogram that has been  
19 identified as the starting subprogram of the program. If more than one  
20 compilation unit contains this flag, any one of them may contain the starting  
21 function.

22 *Fortran has a PROGRAM statement which is used to specify and provide a*  
23 *user-specified name for the main subroutine of a program. C uses the name "main" to*  
24 *identify the main subprogram of a program. Some other languages provide similar or*  
25 *other means to identify the main subprogram of a program. The*  
26 *[DW\\_AT\\_main\\_subprogram](#) attribute may also be used to identify such subprograms*  
27 *(see Section 3.3.1 on [page 78](#)).*

- 28 13. A **DW\_AT\_entry\_pc** attribute whose value is the address of the first  
29 executable instruction of the unit (see Section 2.18 on [page 56](#)).

- 30 14. A **DW\_AT\_str\_offsets** attribute, whose value is of class **stroffsetsptr**. This  
31 attribute points to the header of the compilation unit's contribution to the  
32 `.debug_str_offsets` (or `.debug_str_offsets.dwo`) section. Indirect string  
33 references (using [DW\\_FORM\\_strx](#), [DW\\_FORM\\_strx1](#), [DW\\_FORM\\_strx2](#),  
34 [DW\\_FORM\\_strx3](#) or [DW\\_FORM\\_strx4](#)) within the compilation unit are  
35 interpreted as indices into the array of offsets following that header.

- 1 15. A `DW_AT_addr_base` attribute, whose value is of class `addrptr`. This  
2 attribute points to the beginning of the compilation unit's contribution to the  
3 `.debug_addr` section. Indirect references (using `DW_FORM_addrx`,  
4 `DW_FORM_addrx1`, `DW_FORM_addrx2`, `DW_FORM_addrx3`,  
5 `DW_FORM_addrx4`, `DW_OP_addrx`, `DW_OP_constx`,  
6 `DW_LLE_base_addressx`, `DW_LLE_startx_endx`, `DW_LLE_startx_length`,  
7 `DW_RLE_base_addressx`, `DW_RLE_startx_endx` or `DW_RLE_startx_length`)  
8 within the compilation unit are interpreted as indices relative to this base.
- 9 16. A `DW_AT_rnglists_base` attribute, whose value is of class `rnglistsptr`. This  
10 attribute points to the beginning of the offsets table (immediately following  
11 the header) of the compilation unit's contribution to the `.debug_rnglists`  
12 section. References to range lists (using `DW_FORM_rnglistx`) within the  
13 compilation unit are interpreted relative to this base.
- 14 17. A `DW_AT_loclists_base` attribute, whose value is of class `loclistsptr`. This  
15 attribute points to the beginning of the offsets table (immediately following  
16 the header) of the compilation unit's contribution to the `.debug_loclists`  
17 section. References to location lists (using `DW_FORM_loclistx`) within the  
18 compilation unit are interpreted relative to this base.

19 The base address of a compilation unit is defined as the value of the  
20 `DW_AT_low_pc` attribute, if present; otherwise, it is undefined. If the base  
21 address is undefined, then any DWARF entry or structure defined in terms of the  
22 base address of that compilation unit is not valid.

### 23 3.1.2 Skeleton Compilation Unit Entries

24 When generating a split DWARF object file (see Section 7.3.2 on page 195), the  
25 compilation unit in the `.debug_info` section is a "skeleton" compilation unit with  
26 the tag `DW_TAG_skeleton_unit`, which contains a `DW_AT_dwo_name` attribute  
27 as well as a subset of the attributes of a full or partial compilation unit. In  
28 general, it contains those attributes that are necessary for the consumer to locate  
29 the object file where the split full compilation unit can be found, and for the  
30 consumer to interpret references to addresses in the program.

31 A skeleton compilation unit has no children.

32 A skeleton compilation unit has the following attributes:

- 33 1. A `DW_AT_dwo_name` attribute whose value is a null-terminated string  
34 containing the full or relative path name (relative to the value of the  
35 `DW_AT_comp_dir` attribute, see below) of the object file that contains the full  
36 compilation unit.

## Chapter 3. Program Scope Entries

1 The value in the `dwo_id` field of the unit header for this unit is the same as the  
2 value in the `dwo_id` field of the unit header of the corresponding full  
3 compilation unit (see Section 7.5.1 on page 209).

4 *The means of determining a compilation unit ID does not need to be similar or related*  
5 *to the means of determining a type unit signature. However, it should be suitable for*  
6 *detecting file version skew or other kinds of mismatched files and for looking up a full*  
7 *split unit in a DWARF package file (see Section 7.3.5 on page 198).*

8 A skeleton compilation unit may have additional attributes, which are the same  
9 as for conventional compilation unit entries except as noted, from among the  
10 following:

- 11 2. Either a `DW_AT_low_pc` and `DW_AT_high_pc` pair of attributes or a  
12 `DW_AT_ranges` attribute.
- 13 3. A `DW_AT_stmt_list` attribute.
- 14 4. A `DW_AT_comp_dir` attribute.
- 15 5. A `DW_AT_use_UTF8` attribute.

16 *This attribute applies to strings referred to by the skeleton compilation unit entry*  
17 *itself, and strings in the associated line number information. The representation for*  
18 *strings in the object file referenced by the `DW_AT_dwo_name` attribute is determined*  
19 *by the presence of a `DW_AT_use_UTF8` attribute in the full compilation unit (see*  
20 *Section 3.1.3 on the next page).*

- 21 6. A `DW_AT_str_offsets` attribute, for indirect strings references from the  
22 corresponding split full compilation unit.
- 23 7. A `DW_AT_addr_base` attribute.
- 24 8. A `DW_AT_rnglists_base` attribute, for range list entry references from the  
25 corresponding split full compilation unit.

26 All other attributes of a compilation unit entry (described in Section 3.1.1 on  
27 page 62) are placed in the split full compilation unit (see 3.1.3 on the next page).  
28 The attributes provided by the skeleton compilation unit entry do not need to be  
29 repeated in the full compilation unit entry.

30 *The `DW_AT_addr_base`, `DW_AT_str_offsets`, and `DW_AT_rnglists_base` attributes*  
31 *provide context that may be necessary to interpret the contents of the corresponding split*  
32 *DWARF object file.*

33 *The `DW_AT_base_types` attribute is not defined for a skeleton compilation unit.*

### 3.1.3 Split Full Compilation Unit Entries

A **split full compilation unit** is represented by a debugging information entry with tag `DW_TAG_compile_unit`. It is very similar to a conventional full compilation unit but is logically paired with a specific skeleton compilation unit while being physically separate.

A split full compilation unit may have the following attributes, which are the same as for conventional compilation unit entries except as noted:

1. A `DW_AT_name` attribute.
2. A `DW_AT_language_name` attribute.
3. A `DW_AT_language_version` attribute.
4. A `DW_AT_macros` attribute. The value of this attribute is of class `macptr`, which is an offset relative to the `.debug_macro.dwo` section.
5. A `DW_AT_producer` attribute.
6. A `DW_AT_identifier_case` attribute.
7. A `DW_AT_main_subprogram` attribute.
8. A `DW_AT_entry_pc` attribute.
9. A `DW_AT_use_UTF8` attribute.

*The following attributes are not part of a split full compilation unit entry but instead are inherited (if present) from the corresponding skeleton compilation unit:*

*`DW_AT_low_pc`, `DW_AT_high_pc`, `DW_AT_ranges`, `DW_AT_stmt_list`, `DW_AT_comp_dir`, `DW_AT_str_offsets` and `DW_AT_addr_base`.*

*The `DW_AT_base_types` attribute is not defined for a split full compilation unit.*

### 3.1.4 Type Unit Entries

An object file may contain any number of separate type unit entries, each representing a single complete type definition. Each type unit must be uniquely identified by an 8-byte signature, stored as part of the type unit, which can be used to reference the type definition from debugging information entries in other compilation units and type units.

Conventional and split type units are identical except for the sections in which they are represented (see Section 7.3.2 on page 195 for details). Moreover, the `DW_AT_str_offsets` attribute (see below) is not used in a split type unit.

## Chapter 3. Program Scope Entries

1 A type unit is represented by a debugging information entry with the tag  
2 `DW_TAG_type_unit`. A type unit entry owns debugging information entries that  
3 represent the definition of a single type, plus additional debugging information  
4 entries that may be necessary to include as part of the definition of the type.

5 A type unit entry may have the following attributes:

- 6 1. A `DW_AT_language_name` attribute, whose constant value is an integer  
7 code indicating the source language used to define the type. The set of  
8 language names and their meanings are given in Table 3.1 on page 64.
- 9 2. A `DW_AT_language_version` attribute, whose constant value is an integer  
10 code indicating the source language version as described in Table 3.2 on  
11 page 66.
- 12 3. A `DW_AT_stmt_list` attribute whose value of class `lineptr` points to the line  
13 number information for this type unit.

14 *Because type units do not describe any code, they do not actually need a line number*  
15 *table, but the line number headers contain a list of directories and file names that may*  
16 *be referenced by the `DW_AT_decl_file` attribute of the type or part of its description.*

17 *In an object file with a conventional compilation unit entry, the type unit entries may*  
18 *refer to (share) the line number table used by the compilation unit. In a type unit*  
19 *located in a split compilation unit, the `DW_AT_stmt_list` attribute refers to a*  
20 *“specialized” line number table in the `.debug_line.dwo` section, which contains*  
21 *only the list of directories and file names.*

22 *All type unit entries in a split DWARF object file may (but are not required to) refer*  
23 *to the same specialized line number table.*

- 24 4. A `DW_AT_use_UTF8` attribute, which is a flag whose presence indicates that  
25 all strings referred to by this type unit entry, its children, and its associated  
26 specialized line number table, are represented using the UTF-8  
27 representation.
- 28 5. A `DW_AT_str_offsets` attribute, whose value is of class `stroffsetsptr`. This  
29 attribute points to the header of the type unit’s contribution to the  
30 `.debug_str_offsets` section. Indirect string references (using  
31 `DW_FORM_strx`, `DW_FORM_strx1`, `DW_FORM_strx2`, `DW_FORM_strx3` or  
32 `DW_FORM_strx4`) within the type unit are interpreted as indices into the  
33 array of offsets following that header.

## Chapter 3. Program Scope Entries

1 A type unit entry for a given type T owns a debugging information entry that  
2 represents a defining declaration of type T. If the type is nested within enclosing  
3 types or namespaces, the debugging information entry for T is nested within  
4 debugging information entries describing its containers; otherwise, T is a direct  
5 child of the type unit entry.

6 A type unit entry may also own additional debugging information entries that  
7 represent declarations of additional types that are referenced by type T and have  
8 not themselves been placed in separate type units. Like T, if an additional type U  
9 is nested within enclosing types or namespaces, the debugging information entry  
10 for U is nested within entries describing its containers; otherwise, U is a direct  
11 child of the type unit entry.

12 The containing entries for types T and U are declarations, and the outermost  
13 containing entry for any given type T or U is a direct child of the type unit entry.  
14 The containing entries may be shared among the additional types and between T  
15 and the additional types.

16 *Examples of these kinds of relationships are found in Section E.2.1 on page 398 and*  
17 *Section E.2.3 on page 408.*

18 *Types are not required to be placed in type units. In general, only large types such as*  
19 *structure, class, enumeration, and union types included from header files should be*  
20 *considered for separate type units. Base types and other small types are not usually worth*  
21 *the overhead of placement in separate type units. Types that are unlikely to be replicated,*  
22 *such as those defined in the main source file, are also better left in the main compilation*  
23 *unit.*

## 24 **3.2 Module, Namespace and Importing Entries**

25 *Modules and namespaces provide a means to collect related entities into a single entity*  
26 *and to manage the names of those entities.*

### 27 **3.2.1 Module Entries**

28 *Several languages have the concept of a “module.” A Modula-2 definition module may be*  
29 *represented by a module entry containing a declaration attribute ([DW\\_AT\\_declaration](#)).*  
30 *A Fortran 90 module may also be represented by a module entry (but no declaration*  
31 *attribute is warranted because Fortran has no concept of a corresponding module body).*



## Chapter 3. Program Scope Entries

1 A module is represented by a debugging information entry with the tag  
2 `DW_TAG_module`. Module entries may own other debugging information  
3 entries describing program entities whose declaration scopes end at the end of  
4 the module itself.

5 If the module has a name, the module entry has a `DW_AT_name` attribute whose  
6 value is a null-terminated string containing the module name.

7 The module entry may have either a `DW_AT_low_pc` and `DW_AT_high_pc` pair  
8 of attributes or a `DW_AT_ranges` attribute whose values encode the contiguous  
9 or non-contiguous address ranges, respectively, of the machine instructions  
10 generated for the module initialization code (see Section 2.17 on page 52). It may  
11 also have a `DW_AT_entry_pc` attribute whose value is the address of the first  
12 executable instruction of that initialization code (see Section 2.18 on page 56).

13 If the module has been assigned a priority, it may have a `DW_AT_priority`  
14 attribute. The value of this attribute is a reference to another debugging  
15 information entry describing a variable with a constant value. The value of this  
16 variable is the actual constant value of the module's priority, represented as it  
17 would be on the target architecture.

### 18 3.2.2 Namespace Entries

19 *C++ has the notion of a namespace, which provides a way to implement name hiding, so  
20 that names of unrelated things do not accidentally clash in the global namespace when an  
21 application is linked together.*

22 A namespace is represented by a debugging information entry with the tag  
23 `DW_TAG_namespace`. A namespace extension is represented by a  
24 `DW_TAG_namespace` entry with a `DW_AT_extension` attribute referring to the  
25 previous extension, or if there is no previous extension, to the original  
26 `DW_TAG_namespace` entry. A namespace extension entry does not need to  
27 duplicate information in a previous extension entry of the namespace nor need it  
28 duplicate information in the original namespace entry. (Thus, for a namespace  
29 with a name, a `DW_AT_name` attribute need only be attached directly to the  
30 original `DW_TAG_namespace` entry.)

31 Namespace and namespace extension entries may own other debugging  
32 information entries describing program entities whose declarations occur in the  
33 namespace.

34 A namespace may have a `DW_AT_export_symbols` attribute which is a `flag`  
35 which indicates that all member names defined within the namespace may be  
36 referenced as if they were defined within the containing namespace.



## Chapter 3. Program Scope Entries

1 *This may be used to describe an inline namespace in C++.*

2 If a type, variable, or function declared in a namespace is defined outside of the  
3 body of the namespace declaration, that type, variable, or function definition  
4 entry has a [DW\\_AT\\_specification](#) attribute whose value is a [reference](#) to the  
5 debugging information entry representing the declaration of the type, variable or  
6 function. Type, variable, or function entries with a [DW\\_AT\\_specification](#)  
7 attribute do not need to duplicate information provided by the declaration entry  
8 referenced by the specification attribute.

9 *The C++ global namespace (the namespace referred to by `::f`, for example) is not  
10 explicitly represented in DWARF with a namespace entry (thus mirroring the situation  
11 in C++ source). Global items may be simply declared with no reference to a namespace.*

12 *The C++ compilation unit specific “unnamed namespace” may be represented by a  
13 namespace entry with no name attribute in the original namespace declaration entry  
14 (and therefore no name attribute in any namespace extension entry of this namespace).  
15 C++ states that declarations in the unnamed namespace are implicitly available in the  
16 containing scope; a producer should make this effect explicit with the  
17 [DW\\_AT\\_export\\_symbols](#) attribute, or by using a [DW\\_TAG\\_imported\\_module](#) that is a  
18 sibling of the namespace entry and references it.*

19 *A compiler emitting namespace information may choose to explicitly represent  
20 namespace extensions, or to represent the final namespace declaration of a compilation  
21 unit; this is a quality-of-implementation issue and no specific requirements are given  
22 here. If only the final namespace is represented, it is impossible for a debugger to interpret  
23 using declaration references in exactly the manner defined by the C++ language.*

24 *For C++ namespace examples, see Appendix [D.3](#) on page 331.*

### 25 **3.2.3 Imported (or Renamed) Declaration Entries**

26 *Some languages support the concept of importing into or making accessible in a given  
27 unit certain declarations that occur in a different module or scope. An imported  
28 declaration may sometimes be given another name.*

29 An imported declaration is represented by one or more debugging information  
30 entries with the tag [DW\\_TAG\\_imported\\_declaration](#). When an overloaded entity  
31 is imported, there is one imported declaration entry for each overloading. Each  
32 imported declaration entry has a [DW\\_AT\\_import](#) attribute, whose value is a  
33 [reference](#) to the debugging information entry representing the declaration that is  
34 being imported.

## Chapter 3. Program Scope Entries

1 An imported declaration may also have a `DW_AT_name` attribute whose value is  
2 a null-terminated string containing the name by which the imported entity is to  
3 be known in the context of the imported declaration entry (which may be  
4 different than the name of the entity being imported). If no name is present, then  
5 the name by which the entity is to be known is the same as the name of the entity  
6 being imported.

7 An imported declaration entry with a name attribute may be used as a general  
8 means to rename or provide an alias for an entity, regardless of the context in  
9 which the importing declaration or the imported entity occurs.

10 *A C++ namespace alias may be represented by an imported declaration entry with a*  
11 *name attribute whose value is a null-terminated string containing the alias name and a*  
12 *`DW_AT_import` attribute whose value is a [reference](#) to the applicable original namespace*  
13 *or namespace extension entry.*

14 *A C++ using declaration may be represented by one or more imported declaration entries.*  
15 *When the using declaration refers to an overloaded function, there is one imported*  
16 *declaration entry corresponding to each overloading. Each imported declaration entry*  
17 *has no name attribute but it does have a `DW_AT_import` attribute that refers to the entry*  
18 *for the entity being imported. (C++ provides no means to “rename” an imported entity,*  
19 *other than a namespace).*

20 *A Fortran use statement with an “only list” may be represented by a series of imported*  
21 *declaration entries, one (or more) for each entity that is imported. An entity that is*  
22 *renamed in the importing context may be represented by an imported declaration entry*  
23 *with a name attribute that specifies the new local name.*

### 24 **3.2.4 Imported Module Entries**

25 *Some languages support the concept of importing into or making accessible in a given*  
26 *unit all of the declarations contained within a separate module or namespace.*

27 An imported module declaration is represented by a debugging information  
28 entry with the tag `DW_TAG_imported_module`. An imported module entry  
29 contains a `DW_AT_import` attribute whose value is a [reference](#) to the module or  
30 namespace entry containing the definition and/or declaration entries for the  
31 entities that are to be imported into the context of the imported module entry.

## Chapter 3. Program Scope Entries

1 An imported module declaration may own a set of imported declaration entries,  
2 each of which refers to an entry in the module whose corresponding entity is to  
3 be known in the context of the imported module declaration by a name other  
4 than its name in that module. Any entity in the module that is not renamed in  
5 this way is known in the context of the imported module entry by the same name  
6 as it is declared in the module.

7 *A C++ using directive may be represented by an imported module entry, with a*  
8 *DW\_AT\_import attribute referring to the namespace entry of the appropriate extension*  
9 *of the namespace (which might be the original namespace entry) and no owned entries.*

10 *A Fortran use statement with a “rename list” may be represented by an imported module*  
11 *entry with an import attribute referring to the module and owned entries corresponding*  
12 *to those entities that are renamed as part of being imported.*

13 *A Fortran use statement with neither a “rename list” nor an “only list” may be*  
14 *represented by an imported module entry with an import attribute referring to the*  
15 *module and no owned child entries.*

16 *A use statement with an “only list” is represented by a series of individual imported*  
17 *declaration entries as described in Section 3.2.3 on page 75.*

18 *A Fortran use statement for an entity in a module that is itself imported by a use*  
19 *statement without an explicit mention may be represented by an imported declaration*  
20 *entry that refers to the original debugging information entry. For example, given*

```
module A
integer X, Y, Z
end module

module B
use A
end module

module C
use B, only Q => X
end module
```

21 *the imported declaration entry for Q within module C refers directly to the variable*  
22 *declaration entry for X in module A because there is no explicit representation for X in*  
23 *module B.*

24 *A similar situation arises for a C++ using declaration that imports an entity in terms of*  
25 *a namespace alias. See Appendix D.3 on page 331 for an example.*

### 3.2.5 Imported Unit Entries

The place where a normal or partial compilation unit is imported is represented by a debugging information entry with the tag `DW_TAG_imported_unit`. An imported unit entry contains a `DW_AT_import` attribute whose value is a [reference](#) to the normal or partial compilation unit entry whose declarations logically belong at the place of the imported unit entry.

*An imported unit entry does not necessarily correspond to any entity or construct in the source program. It is merely “glue” used to relate a partial unit, or a compilation unit used as a partial unit, to a place in some other compilation unit.*

## 3.3 Subroutine and Entry Point Entries

The following tags exist to describe debugging information entries for subroutines and entry points:

<code>DW_TAG_subprogram</code>	A subroutine or function
<code>DW_TAG_inlined_subroutine</code>	A particular inlined instance of a subroutine or function
<code>DW_TAG_entry_point</code>	An alternate entry point

### 3.3.1 General Subroutine and Entry Point Information

The subroutine or entry point entry has a `DW_AT_name` attribute whose value is a null-terminated string containing the subroutine or entry point name. It may also have a `DW_AT_linkage_name` attribute as described in [Section 2.22 on page 58](#).

If the name of the subroutine described by an entry with the tag `DW_TAG_subprogram` is visible outside of its containing compilation unit, that entry has a `DW_AT_external` attribute, which is a [flag](#).

*Additional attributes for functions that are members of a class or structure are described in [Section 5.7.8 on page 124](#).*

A subroutine entry may contain a `DW_AT_main_subprogram` attribute which is a [flag](#) whose presence indicates that the subroutine has been identified as the starting function of the program. If more than one subprogram contains this flag, any one of them may be the starting subroutine of the program.

*See also [Section 3.1 on page 61](#)) regarding the related use of this attribute to indicate that a compilation unit contains the main subroutine of a program.*

1 **3.3.1.1 Calling Convention Information**

2 A subroutine entry may contain a `DW_AT_calling_convention` attribute, whose  
3 value is an `integer constant`. The set of calling convention codes for subroutines  
4 is given in Table 3.4.

Table 3.4: Calling convention codes for subroutines

---

`DW_CC_normal`  
`DW_CC_program`  
`DW_CC_nocall`

---

5 If this attribute is not present, or its value is the constant `DW_CC_normal`, then  
6 the subroutine may be safely called by obeying the “standard” calling  
7 conventions of the target architecture. If the value of the calling convention  
8 attribute is the constant `DW_CC_nocall`, the subroutine does not obey standard  
9 calling conventions, and it may not be safe for the debugger to call this  
10 subroutine.

11 *Note that `DW_CC_normal` is also used as a calling convention code for certain types (see*  
12 *Table 5.5 on page 120).*

13 If the semantics of the language of the compilation unit containing the  
14 subroutine entry distinguishes between ordinary subroutines and subroutines  
15 that can serve as the “main program,” that is, subroutines that cannot be called  
16 directly according to the ordinary calling conventions, then the debugging  
17 information entry for such a subroutine may have a calling convention attribute  
18 whose value is the constant `DW_CC_program`.

19 *A common debugger feature is to allow the debugger user to call a subroutine within the*  
20 *subject program. In certain cases, however, the generated code for a subroutine will not*  
21 *obey the standard calling conventions for the target architecture and will therefore not be*  
22 *safe to call from within a debugger.*

23 *The `DW_CC_program` value is intended to support Fortran main programs which in*  
24 *some implementations may not be callable or which must be invoked in a special way. It*  
25 *is not intended as a way of finding the entry address for the program.*

26 **3.3.1.2 Miscellaneous Subprogram Properties**

27 *In C there is a difference between the types of functions declared using function prototype*  
28 *style declarations and those declared using non-prototype declarations.*

## Chapter 3. Program Scope Entries

1 A subroutine entry declared with a function prototype style declaration may  
2 have a `DW_AT_prototyped` attribute, which is a `flag`. The attribute indicates  
3 whether a subroutine entry point corresponds to a function declaration that  
4 includes parameter prototype information.

5 A subprogram entry may have a `DW_AT_elemental` attribute, which is a `flag`.  
6 The attribute indicates whether the subroutine or entry point was declared with  
7 the “elemental” keyword or property.

8 A subprogram entry may have a `DW_AT_pure` attribute, which is a `flag`. The  
9 attribute indicates whether the subroutine was declared with the “pure”  
10 keyword or property.

11 A subprogram entry may have a `DW_AT_recursive` attribute, which is a `flag`. The  
12 attribute indicates whether the subroutine or entry point was declared with the  
13 “recursive” keyword or property.

14 A subprogram entry may have a `DW_AT_noreturn` attribute, which is a `flag`. The  
15 attribute indicates whether the subprogram was declared with the “noreturn”  
16 keyword or property indicating that the subprogram can be called, but will never  
17 return to its caller.

18 *The Fortran language allows the keywords `elemental`, `pure` and `recursive` to be*  
19 *included as part of the declaration of a subroutine; these attributes reflect that usage.*  
20 *These attributes are not relevant for languages that do not support similar keywords or*  
21 *syntax. In particular, the `DW_AT_recursive` attribute is neither needed nor appropriate*  
22 *in languages such as C where functions support recursion by default.*

### 23 3.3.1.3 Call Site-Related Attributes

24 *While subprogram attributes in the previous section provide information about the*  
25 *subprogram and its entry point(s) as a whole, the following attributes provide summary*  
26 *information about the calls that occur within a subprogram.*

27 A subroutine entry may have `DW_AT_call_all_tail_calls`, `DW_AT_call_all_calls`  
28 and/or `DW_AT_call_all_source_calls` attributes, each of which is a `flag`. These  
29 flags indicate the completeness of the call site information provided by call site  
30 entries (see Section 3.4.1 on page 92) within the subprogram.

31 The `DW_AT_call_all_tail_calls` attribute indicates that every tail call that occurs  
32 in the code for the subprogram is described by a `DW_TAG_call_site` entry. (There  
33 may or may not be other non-tail calls to some of the same target subprograms.)

34 The `DW_AT_call_all_calls` attribute indicates that every non-inlined call (either a  
35 tail call or a normal call) that occurs in the code for the subprogram is described  
36 by a `DW_TAG_call_site` entry.



## Chapter 3. Program Scope Entries

1 The `DW_AT_call_all_source_calls` attribute indicates that every call that occurs in  
2 the code for the subprogram, including every call inlined into it, is described by  
3 either a `DW_TAG_call_site` entry or a `DW_TAG_inlined_subroutine` entry;  
4 further, any call that is optimized out is nonetheless also described using a  
5 `DW_TAG_call_site` entry that has neither a `DW_AT_call_pc` nor  
6 `DW_AT_call_return_pc` attribute.

7 *The `DW_AT_call_all_source_calls` attribute is intended for debugging information*  
8 *format consumers that analyze call graphs.*

9 If the the `DW_AT_call_all_source_calls` attribute is present then the  
10 `DW_AT_call_all_calls` and `DW_AT_call_all_tail_calls` attributes are also  
11 implicitly present. Similarly, if the `DW_AT_call_all_calls` attribute is present then  
12 the `DW_AT_call_all_tail_calls` attribute is implicitly present.

### 13 3.3.2 Subroutine and Entry Point Return Types

14 If the subroutine or entry point is a function that returns a value, then its  
15 debugging information entry has a `DW_AT_type` attribute to denote the type  
16 returned by that function.

17 *Debugging information entries for C void functions should not have an attribute for the*  
18 *return type.*

19 *Debugging information entries for declarations of C++ member functions with an `auto`*  
20 *return type specifier should use an unspecified type entry (see Section 5.2 on page 112).*  
21 *The debugging information entry for the corresponding definition should provide the*  
22 *deduced return type. This practice causes the description of the containing class to be*  
23 *consistent across compilation units, allowing the class declaration to be placed into a*  
24 *separate type unit if desired.*

### 25 3.3.3 Subroutine and Entry Point Locations

26 A subroutine entry may have either a `DW_AT_low_pc` and `DW_AT_high_pc`  
27 pair of attributes or a `DW_AT_ranges` attribute whose values encode the  
28 contiguous or non-contiguous address ranges, respectively, of the machine  
29 instructions generated for the subroutine (see Section 2.17 on page 52).

30 A subroutine entry may also have a `DW_AT_entry_pc` attribute whose value is  
31 the address of the first executable instruction of the subroutine (see Section 2.18  
32 on page 56).

33 An entry point has a `DW_AT_low_pc` attribute whose value is the relocated  
34 address of the first machine instruction generated for the entry point.

1 Subroutines and entry points may also have a `DW_AT_address_class` attribute, if  
2 appropriate, to specify the addressing mode to be used in calling that subroutine.

3 A subroutine entry representing a subroutine declaration that is not also a  
4 definition does not have code address or range attributes.

### 5 **3.3.4 Declarations Owned by Subroutines and Entry Points**

6 The declarations enclosed by a subroutine or entry point are represented by  
7 debugging information entries that are owned by the subroutine or entry point  
8 entry. Entries representing the formal parameters of the subroutine or entry point  
9 appear in the same order as the corresponding declarations in the source  
10 program.

11 *There is no ordering requirement for entries for declarations other than formal*  
12 *parameters. The formal parameter entries may be interspersed with other entries used by*  
13 *formal parameter entries, such as type entries.*

14 The unspecified (sometimes called “varying”) parameters of a subroutine  
15 parameter list are represented by a debugging information entry with the tag  
16 `DW_TAG_unspecified_parameters`.

17 The entry for a subroutine that includes a Fortran `common block` has a child  
18 entry with the tag `DW_TAG_common_inclusion`. The common inclusion entry  
19 has a `DW_AT_common_reference` attribute whose value is a `reference` to the  
20 debugging information entry for the common block being included (see Section  
21 [4.2 on page 102](#)).

### 22 **3.3.5 Low-Level Information**

23 A subroutine or entry point entry may have a `DW_AT_return_addr` attribute,  
24 whose value is a location description. The location specified is the place where  
25 the return address for the subroutine or entry point is stored.

26 A subroutine or entry point entry may also have a `DW_AT_frame_base` attribute,  
27 whose value is a location description that describes the “frame base” for the  
28 subroutine or entry point. If the location description is a simple register location  
29 description, the given register contains the frame base address. If the location  
30 description is a DWARF expression, the result of evaluating that expression is the  
31 frame base address. Finally, for a location list, this interpretation applies to each  
32 location description contained in the list of location list entries.

33 *The use of one of the `DW_OP_reg<n>` operations in this context is equivalent to using*  
34 *`DW_OP_breg<n>(0)` but more compact. However, these are not equivalent in general.*



## Chapter 3. Program Scope Entries

1     *The frame base for a subprogram is typically an address relative to the first unit of storage*  
2     *allocated for the subprogram's stack frame. The [DW\\_AT\\_frame\\_base](#) attribute can be*  
3     *used in several ways:*

- 4     1. *In subprograms that need location lists to locate local variables, the*  
5         *[DW\\_AT\\_frame\\_base](#) can hold the needed location list, while all variables' location*  
6         *descriptions can be simpler ones involving the frame base.*
- 7     2. *It can be used in resolving "up-level" addressing within nested routines. (See also*  
8         *[DW\\_AT\\_static\\_link](#), below)*

9     *Some languages support nested subroutines. In such languages, it is possible to reference*  
10    *the local variables of an outer subroutine from within an inner subroutine. The*  
11    *[DW\\_AT\\_static\\_link](#) and [DW\\_AT\\_frame\\_base](#) attributes allow debuggers to support this*  
12    *same kind of referencing.*

13    If a subroutine or entry point is nested, it may have a [DW\\_AT\\_static\\_link](#)  
14    attribute, whose value is a location description that computes the frame base of  
15    the relevant instance of the subroutine that immediately encloses the subroutine  
16    or entry point.

17    In the context of supporting nested subroutines, the [DW\\_AT\\_frame\\_base](#)  
18    attribute value obeys the following constraints:

- 19    1. It computes a value that does not change during the life of the subprogram,  
20       and
- 21    2. The computed value is unique among instances of the same subroutine.

22         *For typical [DW\\_AT\\_frame\\_base](#) use, this means that a recursive subroutine's stack*  
23         *frame must have non-zero size.*

24    *If a debugger is attempting to resolve an up-level reference to a variable, it uses the*  
25    *nesting structure of DWARF to determine which subroutine is the lexical parent and the*  
26    *[DW\\_AT\\_static\\_link](#) value to identify the appropriate active frame of the parent. It can*  
27    *then attempt to find the reference within the context of the parent.*

### 28    **3.3.6 Types Thrown by Exceptions**

29    *In C++ a subroutine may declare a set of types which it may validly throw.*

30    If a subroutine explicitly declares that it may throw an exception of one or more  
31    types, each such type is represented by a debugging information entry with the  
32    tag [DW\\_TAG\\_thrown\\_type](#). Each such entry is a child of the entry representing  
33    the subroutine that may throw this type. Each thrown type entry contains a  
34    [DW\\_AT\\_type](#) attribute, whose value is a [reference](#) to an entry describing the type  
35    of the exception that may be thrown.

### 3.3.7 Function Template Instantiations

*In C++, a function template is a generic definition of a function that is instantiated differently for calls with values of different types. DWARF does not represent the generic template definition, but does represent each instantiation.*

A function template instantiation is represented by a debugging information entry with the tag `DW_TAG_subprogram`. With the following exceptions, such an entry will contain the same attributes and will have the same types of child entries as would an entry for a subroutine defined explicitly using the instantiation types and values. The exceptions are:

1. Template parameters are described and referenced as specified in Section 2.23 on page 58.
2. If the compiler has generated a separate compilation unit to hold the template instantiation and that compilation unit has a different name from the compilation unit containing the template definition, the name attribute for the debugging information entry representing that compilation unit is empty or omitted.
3. If the subprogram entry representing the template instantiation or any of its child entries contain declaration coordinate attributes, those attributes refer to the source for the template definition, not to any source generated artificially by the compiler for this instantiation.

### 3.3.8 Inlinable and Inlined Subroutines

A declaration or a definition of an inlinable subroutine is represented by a debugging information entry with the tag `DW_TAG_subprogram`. The entry for a subroutine that is explicitly declared to be available for inline expansion or that was expanded inline implicitly by the compiler has a `DW_AT_inline` attribute whose value is an `integer constant`. The set of values for the `DW_AT_inline` attribute is given in Table 3.5 on the next page.

*In C++, a function or a constructor declared with `constexpr` is implicitly declared inline. The abstract instance (see Section 3.3.8.1 on the following page) is represented by a debugging information entry with the tag `DW_TAG_subprogram`. Such an entry has a `DW_AT_inline` attribute whose value is `DW_INL_inlined`.*

Table 3.5: Inline codes

Name	Meaning
<code>DW_INL_not_inlined</code>	Not declared inline nor inlined by the compiler (equivalent to the absence of the containing <code>DW_AT_inline</code> attribute)
<code>DW_INL_inlined</code>	Not declared inline but inlined by the compiler
<code>DW_INL_declared_not_inlined</code>	Declared inline but not inlined by the compiler
<code>DW_INL_declared_inlined</code>	Declared inline and inlined by the compiler

### 3.3.8.1 Abstract Instances

Any subroutine entry that contains a `DW_AT_inline` attribute whose value is other than `DW_INL_not_inlined` is known as an **abstract instance root**. Any debugging information entry that is owned (either directly or indirectly) by an abstract instance root is known as an **abstract instance entry**. Any set of abstract instance entries that are all children (either directly or indirectly) of some abstract instance root, together with the root itself, is known as an **abstract instance tree**. However, in the case where an abstract instance tree is nested within another abstract instance tree, the entries in the nested abstract instance tree are not considered to be entries in the outer abstract instance tree.

Each abstract instance root is either part of a larger tree (which gives a context for the root) or uses `DW_AT_specification` to refer to the declaration in context.

*For example, in C++ the context might be a namespace declaration or a class declaration.*

*Abstract instance trees are defined so that no entry is part of more than one abstract instance tree.*

Attributes and children in an abstract instance are shared by all concrete instances (see Section 3.3.8.2 on the next page).

A debugging information entry that is a member of an abstract instance tree may not contain any attributes which describe aspects of the subroutine which vary between distinct inlined expansions or distinct out-of-line expansions.

*For example, the `DW_AT_low_pc`, `DW_AT_high_pc`, `DW_AT_ranges`, `DW_AT_entry_pc`, `DW_AT_location`, `DW_AT_return_addr` and `DW_AT_start_scope` attributes typically should be omitted; however, this list is not exhaustive.*

## Chapter 3. Program Scope Entries

1 *It would not make sense normally to put these attributes into abstract instance entries*  
2 *since such entries do not represent actual (concrete) instances and thus do not actually*  
3 *exist at run-time. However, see Appendix D.7.3 on page 351 for a contrary example.*

4 The rules for the relative location of entries belonging to abstract instance trees  
5 are exactly the same as for other similar types of entries that are not abstract.  
6 Specifically, the rule that requires that an entry representing a declaration be a  
7 direct child of the entry representing the scope of the declaration applies equally  
8 to both abstract and non-abstract entries. Also, the ordering rules for formal  
9 parameter entries, member entries, and so on, all apply regardless of whether or  
10 not a given entry is abstract.

### 11 3.3.8.2 Concrete Instances

12 Each inline expansion of a subroutine is represented by a debugging information  
13 entry with the tag `DW_TAG_inlined_subroutine`. Each such entry is a direct  
14 child of the entry that represents the scope within which the inlining occurs.

15 Each inlined subroutine entry may have either a `DW_AT_low_pc` and  
16 `DW_AT_high_pc` pair of attributes or a `DW_AT_ranges` attribute whose values  
17 encode the contiguous or non-contiguous address ranges, respectively, of the  
18 machine instructions generated for the inlined subroutine (see Section 2.17  
19 following). An inlined subroutine entry may also contain a `DW_AT_entry_pc`  
20 attribute, representing the first executable instruction of the inline expansion (see  
21 Section 2.18 on page 56).

22 An inlined subroutine entry may also have `DW_AT_call_file`, `DW_AT_call_line`  
23 and `DW_AT_call_column` attributes, each of whose value is an *integer constant*.  
24 These attributes represent the source file, source line number, and source column  
25 number, respectively, of the first character of the statement or expression that  
26 caused the inline expansion. The call file, call line, and call column attributes are  
27 interpreted in the same way as the declaration file, declaration line, and  
28 declaration column attributes, respectively (see Section 2.14 on page 50).

29 *The call file, call line and call column coordinates do not describe the coordinates of the*  
30 *subroutine declaration that was inlined, rather they describe the coordinates of the call.*

31 An inlined subroutine entry may have a `DW_AT_const_expr` attribute, which is a  
32 *flag* whose presence indicates that the subroutine has been evaluated as a  
33 compile-time constant. Such an entry may also have a `DW_AT_const_value`  
34 attribute, whose value may be of any form that is appropriate for the  
35 representation of the subroutine's return value. The value of this attribute is the  
36 actual return value of the subroutine, represented as it would be on the target  
37 architecture.

## Chapter 3. Program Scope Entries

1 *In C++, if a function or a constructor declared with `constexpr` is called with constant*  
2 *expressions, then the corresponding concrete inlined instance has a `DW_AT_const_expr`*  
3 *attribute, as well as a `DW_AT_const_value` attribute whose value represents the actual*  
4 *return value of the concrete inlined instance.*

5 Any debugging information entry that is owned (either directly or indirectly) by  
6 a debugging information entry with the tag `DW_TAG_inlined_subroutine` is  
7 referred to as a “concrete inlined instance entry.” Any entry that has the tag  
8 `DW_TAG_inlined_subroutine` is known as a “concrete inlined instance root.”

9 Any set of concrete inlined instance entries that are all children (either directly or  
10 indirectly) of some concrete inlined instance root, together with the root itself, is  
11 known as a “concrete inlined instance tree.” However, in the case where a  
12 concrete inlined instance tree is nested within another concrete instance tree, the  
13 entries in the nested concrete inline instance tree are not considered to be entries  
14 in the outer concrete instance tree.

15 *Concrete inlined instance trees are defined so that no entry is part of more than one*  
16 *concrete inlined instance tree. This simplifies later descriptions.*

17 Each concrete inlined instance tree is uniquely associated with one (and only  
18 one) abstract instance tree.

19 *Note, however, that the reverse is not true. Any given abstract instance tree may be*  
20 *associated with several different concrete inlined instance trees, or may even be associated*  
21 *with zero concrete inlined instance trees.*

22 Concrete inlined instance entries may omit attributes that are not specific to the  
23 concrete instance (but present in the abstract instance) and need include only  
24 attributes that are specific to the concrete instance (but omitted in the abstract  
25 instance). In place of these omitted attributes, each concrete inlined instance  
26 entry has a `DW_AT_abstract_origin` attribute that may be used to obtain the  
27 missing information (indirectly) from the associated abstract instance entry. The  
28 value of the abstract origin attribute is a reference to the associated abstract  
29 instance entry.

30 If an entry within a concrete inlined instance tree contains attributes describing  
31 the `declaration coordinates` of that entry, then those attributes refer to the file, line  
32 and column of the original declaration of the subroutine, not to the point at  
33 which it was inlined. As a consequence, they may usually be omitted from any  
34 entry that has an abstract origin attribute.

## Chapter 3. Program Scope Entries

1 For each pair of entries that are associated via a [DW\\_AT\\_abstract\\_origin](#)  
2 attribute, both members of the pair have the same tag. So, for example, an entry  
3 with the tag [DW\\_TAG\\_variable](#) can only be associated with another entry that  
4 also has the tag [DW\\_TAG\\_variable](#). The only exception to this rule is that the  
5 root of a concrete instance tree (which must always have the tag  
6 [DW\\_TAG\\_inlined\\_subroutine](#)) can only be associated with the root of its  
7 associated abstract instance tree (which must have the tag  
8 [DW\\_TAG\\_subprogram](#)).

9 In general, the structure and content of any given concrete inlined instance tree  
10 will be closely analogous to the structure and content of its associated abstract  
11 instance tree. There are a few exceptions:

- 12 1. An entry in the concrete instance tree may be omitted if it contains only a  
13 [DW\\_AT\\_abstract\\_origin](#) attribute and either has no children, or its children  
14 are omitted. Such entries would provide no useful information. In C-like  
15 languages, such entries frequently include types, including structure, union,  
16 class, and interface types; and members of types. If any entry within a  
17 concrete inlined instance tree needs to refer to an entity declared within the  
18 scope of the relevant inlined subroutine and for which no concrete instance  
19 entry exists, the reference refers to the abstract instance entry.
- 20 2. Entries in the concrete instance tree which are associated with entries in the  
21 abstract instance tree such that neither has a [DW\\_AT\\_name](#) attribute, and  
22 neither is referenced by any other debugging information entry, may be  
23 omitted. This may happen for debugging information entries in the abstract  
24 instance trees that became unnecessary in the concrete instance tree because  
25 of additional information available there. For example, an anonymous  
26 variable might have been created and described in the abstract instance tree,  
27 but because of the actual parameters for a particular inlined expansion, it  
28 could be described as a constant value without the need for that separate  
29 debugging information entry.
- 30 3. A concrete instance tree may contain entries which do not correspond to  
31 entries in the abstract instance tree to describe new entities that are specific to  
32 a particular inlined expansion. In that case, they will not have associated  
33 entries in the abstract instance tree, do not contain [DW\\_AT\\_abstract\\_origin](#)  
34 attributes, and must contain all their own attributes directly. This allows an  
35 abstract instance tree to omit debugging information entries for anonymous  
36 entities that are unlikely to be needed in most inlined expansions. In any  
37 expansion which deviates from that expectation, the entries can be described  
38 in its concrete inlined instance tree.



1     **3.3.8.3 Out-of-Line Instances of Inlined Subroutines**

2     Under some conditions, compilers may need to generate concrete executable  
3     instances of inlined subroutines other than at points where those subroutines are  
4     actually called. Such concrete instances of inlined subroutines are referred to as  
5     “concrete out-of-line instances.”

6     *In C++, for example, taking the address of a function declared to be inline can necessitate*  
7     *the generation of a concrete out-of-line instance of the given function.*

8     The DWARF representation of a concrete out-of-line instance of an inlined  
9     subroutine is essentially the same as for a concrete inlined instance of that  
10    subroutine (as described in the preceding section). The representation of such a  
11    concrete out-of-line instance makes use of **DW\_AT\_abstract\_origin** attributes in  
12    exactly the same way as they are used for a concrete inlined instance (that is, as  
13    references to corresponding entries within the associated abstract instance tree).

14    The differences between the DWARF representation of a concrete out-of-line  
15    instance of a given subroutine and the representation of a concrete inlined  
16    instance of that same subroutine are as follows:

- 17    1. The root entry for a concrete out-of-line instance of a given inlined subroutine  
18    has the same tag as does its associated (abstract) inlined subroutine entry  
19    (that is, tag **DW\_TAG\_subprogram** rather than  
20    **DW\_TAG\_inlined\_subroutine**).
- 21    2. The root entry for a concrete out-of-line instance tree is normally owned by  
22    the same parent entry that also owns the root entry of the associated abstract  
23    instance. However, it is not required that the abstract and out-of-line instance  
24    trees be owned by the same parent entry.

25    **3.3.8.4 Nested Inlined Subroutines**

26    Some languages and compilers may permit the logical nesting of a subroutine  
27    within another subroutine, and may permit either the outer or the nested  
28    subroutine, or both, to be inlined.

29    For a non-inlined subroutine nested within an inlined subroutine, the nested  
30    subroutine is described normally in both the abstract and concrete inlined  
31    instance trees for the outer subroutine. All rules pertaining to the abstract and  
32    concrete instance trees for the outer subroutine apply also to the abstract and  
33    concrete instance entries for the nested subroutine.

## Chapter 3. Program Scope Entries

1 For an inlined subroutine nested within another inlined subroutine, the  
2 following rules apply to their abstract and concrete instance trees:

- 3 1. The abstract instance tree for the nested subroutine is described within the  
4 abstract instance tree for the outer subroutine according to the rules in  
5 Section [3.3.8.1 on page 85](#), and without regard to the fact that it is within an  
6 outer abstract instance tree.
- 7 2. Any abstract instance tree for a nested subroutine is always omitted within  
8 the concrete instance tree for an outer subroutine.
- 9 3. A concrete instance tree for a nested subroutine is always omitted within the  
10 abstract instance tree for an outer subroutine.
- 11 4. The concrete instance tree for any inlined or out-of-line expansion of the  
12 nested subroutine is described within a concrete instance tree for the outer  
13 subroutine according to the rules in Sections [3.3.8.2 on page 86](#) or [3.3.8.3](#)  
14 following, respectively, and without regard to the fact that it is within an  
15 outer concrete instance tree.

16 See *Appendix D.7 on page 347* for discussion and examples.

### 17 3.3.9 Trampolines

18 *A trampoline is a compiler-generated subroutine that serves as an intermediary in*  
19 *making a call to another subroutine. It may adjust parameters and/or the result (if any)*  
20 *as appropriate to the combined calling and called execution contexts.*

21 A trampoline is represented by a debugging information entry with the tag  
22 [DW\\_TAG\\_subprogram](#) or [DW\\_TAG\\_inlined\\_subroutine](#) that has a  
23 [DW\\_AT\\_trampoline](#) attribute. The value of that attribute indicates the target  
24 subroutine of the trampoline, that is, the subroutine to which the trampoline  
25 passes control. (A trampoline entry may but need not also have a  
26 [DW\\_AT\\_artificial](#) attribute.)

27 The value of the trampoline attribute may be represented using any of the  
28 following forms:

- 29 • If the value is of class [reference](#), then the value specifies the debugging  
30 information entry of the target subprogram.
- 31 • If the value is of class [address](#), then the value is the relocated address of the  
32 target subprogram.



## Chapter 3. Program Scope Entries

- 1       • If the value is of class `string`, then the value is the (possibly mangled) name  
2       of the target subprogram.
- 3       • If the value is of class `flag`, then the value true indicates that the containing  
4       subroutine is a trampoline but that the target subroutine is not known.

5       The target subprogram may itself be a trampoline. (A sequence of trampolines  
6       necessarily ends with a non-trampoline subprogram.)

7       *In C++, trampolines may be used to implement derived virtual member functions; such  
8       trampolines typically adjust the implicit `this` parameter in the course of passing control.  
9       Other languages and environments may use trampolines in a manner sometimes known  
10      as transfer functions or transfer vectors.*

11      *Trampolines may sometimes pass control to the target subprogram using a branch or  
12      jump instruction instead of a call instruction, thereby leaving no trace of their existence  
13      in the subsequent execution context.*

14      *This attribute helps make it feasible for a debugger to arrange that stepping into a  
15      trampoline or setting a breakpoint in a trampoline will result in stepping into or setting  
16      the breakpoint in the target subroutine instead. This helps to hide the compiler generated  
17      subprogram from the user.*

### 18      **3.4 Call Site Entries and Parameters**

19      *A call site entry describes a call from one subprogram to another in the source program.  
20      It provides information about the actual parameters of the call so that they may be more  
21      easily accessed by a debugger. When used together with call frame information (see  
22      Section 6.4 on page 178), call site entries can be useful for computing the value of an  
23      actual parameter passed by a caller, even when the location description for the callee's  
24      corresponding formal parameter does not provide a current location for the formal  
25      parameter.*

26      *The DWARF expression for computing the value of an actual parameter at a call site may  
27      refer to registers or memory locations. The expression assumes these contain the values  
28      they would have at the point where the call is executed. After the called subprogram has  
29      been entered, these registers and memory locations might have been modified. In order to  
30      recover the values that existed at the point of the call (to allow evaluation of the DWARF  
31      expression for the actual parameter), a debugger may virtually unwind the subprogram  
32      activation (see Section 6.4 on page 178). Any register or memory location that cannot be  
33      recovered is referred to as "clobbered by the call."*

## Chapter 3. Program Scope Entries

1 A source call can be compiled into different types of machine code:

- 2 • A *normal call* uses a call-like instruction which transfers control to the start  
3 of some subprogram and preserves the call site location for use by the  
4 callee.
- 5 • A *tail call* uses a jump-like instruction which transfers control to the start of  
6 some subprogram, but there is no call site location address to preserve (and  
7 thus none is available using the virtual unwind information).
- 8 • A *tail recursion call* is a call to the current subroutine which is compiled as a  
9 jump to the current subroutine.
- 10 • An *inline (or inlined) call* is a call to an inlined subprogram, where at least  
11 one instruction has the location of the inlined subprogram or any of its  
12 blocks or inlined subprograms.

13 There are also different types of “optimized out” calls:

- 14 • An *optimized out (normal) call* is a call that is in unreachable code that has  
15 not been emitted (such as, for example, the call to `foo` in `if (0) foo();`).
- 16 • An *optimized out inline call* is a call to an inlined subprogram which either  
17 did not expand to any instructions or only parts of instructions belong to it  
18 and for debug information purposes those instructions are given a location  
19 in the caller.

20 [DW\\_TAG\\_call\\_site](#) entries describe normal and tail calls but not tail recursion  
21 calls, while [DW\\_TAG\\_inlined\\_subroutine](#) entries describe inlined calls (see  
22 Section 3.3.8 on page 84). Call site entries cannot fully describe tail recursion or  
23 optimized out calls.

24 *For optimized out calls there is no code address to use for [DW\\_AT\\_call\\_return\\_pc](#) or*  
25 *[DW\\_AT\\_call\\_pc](#) attributes; however, the fact that the source code makes a call to a certain*  
26 *function at a specific source code location and whether some of the arguments have*  
27 *constant values can be useful for certain consumers.*

### 28 3.4.1 Call Site Entries

29 A call site is represented by a debugging information entry with the tag  
30 [DW\\_TAG\\_call\\_site](#). The entry for a call site is owned by the innermost  
31 debugging information entry representing the scope within which the call is  
32 present in the source program.

## Chapter 3. Program Scope Entries

1 *A scope entry (for example, a lexical block) that would not otherwise be present in the*  
2 *debugging information of a subroutine need not be introduced solely to represent the*  
3 *immediately containing scope of a call.*

4 The call site entry may have a `DW_AT_call_return_pc` attribute which is the  
5 return address after the call. The value of this attribute corresponds to the return  
6 address computed by call frame information in the called subprogram (see  
7 Section 7.24 on page 251).

8 *On many architectures the return address is the address immediately following the call*  
9 *instruction, but on architectures with delay slots it might be an address after the delay*  
10 *slot of the call.*

11 The call site entry may have a `DW_AT_call_pc` attribute which is the address of  
12 the call-like instruction for a normal call or the jump-like instruction for a tail call.

13 If the call site entry corresponds to a tail call, it has the `DW_AT_call_tail_call`  
14 attribute, which is a *flag*.

15 The call site entry may have a `DW_AT_call_origin` attribute which is a *reference*.  
16 For direct calls or jumps where the called subprogram is known it is a reference  
17 to the called subprogram's debugging information entry. For indirect calls it may  
18 be a reference to a `DW_TAG_variable`, `DW_TAG_formal_parameter` or  
19 `DW_TAG_member` entry representing the subroutine pointer that is called.

20 The call site may have a `DW_AT_call_target` attribute which is a DWARF  
21 expression. For indirect calls or jumps where it is unknown at compile time  
22 which subprogram will be called the expression computes the address of the  
23 subprogram that will be called.

24 *The DWARF expression should not use register or memory locations that might be*  
25 *clobbered by the call.*

26 The call site entry may have a `DW_AT_call_target_clobbered` attribute which is a  
27 DWARF expression. For indirect calls or jumps where the address is not  
28 computable without use of registers or memory locations that might be  
29 clobbered by the call the `DW_AT_call_target_clobbered` attribute is used instead  
30 of the `DW_AT_call_target` attribute.

31 *The expression of a call target clobbered attribute may only be valid at the time the call or*  
32 *call-like transfer of control is executed.*

33 The call site entry may have a `DW_AT_type` attribute referencing a debugging  
34 information entry for the type of the called function.

35 *When `DW_AT_call_origin` is present, `DW_AT_type` is usually omitted.*

1 The call site entry may have `DW_AT_call_file`, `DW_AT_call_line` and  
2 `DW_AT_call_column` attributes, each of whose value is an integer constant.  
3 These attributes represent the source file, source line number, and source column  
4 number, respectively, of the first character of the call statement or expression.  
5 The call file, call line, and call column attributes are interpreted in the same way  
6 as the declaration file, declaration line, and declaration column attributes,  
7 respectively (see Section 2.14 on page 50).

8 *The call file, call line and call column coordinates do not describe the coordinates of the*  
9 *subroutine declaration that was called, rather they describe the coordinates of the call.*

### 10 3.4.2 Call Site Parameters

11 The call site entry may own `DW_TAG_call_site_parameter` debugging  
12 information entries representing the parameters passed to the call. Call site  
13 parameter entries occur in the same order as the corresponding parameters in the  
14 source. Each such entry has a `DW_AT_location` attribute which is a location  
15 description. This location description describes where the parameter is passed  
16 (usually either some register, or a memory location expressible as the contents of  
17 the stack register plus some offset).

18 Each `DW_TAG_call_site_parameter` entry may have a `DW_AT_call_value`  
19 attribute which is a DWARF expression which when evaluated yields the value  
20 of the parameter at the time of the call.

21 *If it is not possible to avoid registers or memory locations that might be clobbered by the*  
22 *call in the expression, then the `DW_AT_call_value` attribute should not be provided. The*  
23 *reason for the restriction is that the value of the parameter may be needed in the midst of*  
24 *the callee, where the call clobbered registers or memory might be already clobbered, and if*  
25 *the consumer is not assured by the producer it can safely use those values, the consumer*  
26 *can not safely use the values at all.*

27 For parameters passed by reference, where the code passes a pointer to a location  
28 which contains the parameter, or for reference type parameters, the  
29 `DW_TAG_call_site_parameter` entry may also have a `DW_AT_call_data_location`  
30 attribute whose value is a location description and a `DW_AT_call_data_value`  
31 attribute whose value is a DWARF expression. The `DW_AT_call_data_location`  
32 attribute describes where the referenced value lives during the call. If it is just  
33 `DW_OP_push_object_address`, it may be left out. The `DW_AT_call_data_value`  
34 attribute describes the value in that location. The expression should not use  
35 registers or memory locations that might be clobbered by the call, as it might be  
36 evaluated after virtually unwinding from the called function back to the caller.

1 Each call site parameter entry may also have a `DW_AT_call_parameter` attribute  
2 which contains a reference to a `DW_TAG_formal_parameter` entry, `DW_AT_type`  
3 attribute referencing the type of the parameter or `DW_AT_name` attribute  
4 describing the parameter's name.

5 *Examples using call site entries and related attributes are found in Appendix D.15 on*  
6 *page 371.*

### 7 **3.5 Lexical Block Entries**

8 *A lexical block is a bracketed sequence of source statements that may contain any number*  
9 *of declarations. In some languages (including C and C++), blocks can be nested within*  
10 *other blocks to any depth.*

11 A lexical block is represented by a debugging information entry with the tag  
12 `DW_TAG_lexical_block`.

13 The lexical block entry may have either a `DW_AT_low_pc` and `DW_AT_high_pc`  
14 pair of attributes or a `DW_AT_ranges` attribute whose values encode the  
15 contiguous or non-contiguous address ranges, respectively, of the machine  
16 instructions generated for the lexical block (see Section 2.17 on page 52).

17 A lexical block entry may also have a `DW_AT_entry_pc` attribute whose value is  
18 the address of the first executable instruction of the lexical block (see Section 2.18  
19 on page 56).

20 If a name has been given to the lexical block in the source program, then the  
21 corresponding lexical block entry has a `DW_AT_name` attribute whose value is a  
22 null-terminated string containing the name of the lexical block.

23 *This is not the same as a C or C++ label (see Section 3.6 on the following page).*

24 The lexical block entry owns debugging information entries that describe the  
25 declarations within that lexical block. There is one such debugging information  
26 entry for each local declaration of an identifier or inner lexical block.

## 3.6 Label Entries

*A label is a way of identifying a source location. A labeled statement is usually the target of one or more “go to” statements.*

A label is represented by a debugging information entry with the tag `DW_TAG_label`. The entry for a label is owned by the debugging information entry representing the scope within which the name of the label could be legally referenced within the source program.

The label entry has a `DW_AT_low_pc` attribute whose value is the address of the first executable instruction for the location identified by the label in the source program. The label entry also has a `DW_AT_name` attribute whose value is a null-terminated string containing the name of the label.

## 3.7 With Statement Entries

*Both Pascal and Modula-2 support the concept of a “with” statement. The with statement specifies a sequence of executable statements within which the fields of a record variable may be referenced, unqualified by the name of the record variable.*

A with statement is represented by a debugging information entry with the tag `DW_TAG_with_stmt`.

A with statement entry may have either a `DW_AT_low_pc` and `DW_AT_high_pc` pair of attributes or a `DW_AT_ranges` attribute whose values encode the contiguous or non-contiguous address ranges, respectively, of the machine instructions generated for the with statement (see Section 2.17 on page 52).

A with statement entry may also have a `DW_AT_entry_pc` attribute whose value is the address of the first executable instruction of the with statement (see Section 2.18 on page 56).

The with statement entry has a `DW_AT_type` attribute, denoting the type of record whose fields may be referenced without full qualification within the body of the statement. It also has a `DW_AT_location` attribute, describing how to find the base address of the record object referenced within the body of the with statement.

## 3.8 Try and Catch Block Entries

*In C++, a [lexical block](#) may be designated as a “catch block.” A catch block is an exception handler that handles exceptions thrown by an immediately preceding “try block.” A catch block designates the type of the exception that it can handle.*

A try block is represented by a debugging information entry with the tag [DW\\_TAG\\_try\\_block](#). A catch block is represented by a debugging information entry with the tag [DW\\_TAG\\_catch\\_block](#).

Both try and catch block entries may have either a [DW\\_AT\\_low\\_pc](#) and [DW\\_AT\\_high\\_pc](#) pair of attributes or a [DW\\_AT\\_ranges](#) attribute whose values encode the contiguous or non-contiguous address ranges, respectively, of the machine instructions generated for the block (see [Section 2.17 on page 52](#)).

A try or catch block entry may also have a [DW\\_AT\\_entry\\_pc](#) attribute whose value is the address of the first executable instruction of the try or catch block (see [Section 2.18 on page 56](#)).

Catch block entries have at least one child entry, an entry representing the type of exception accepted by that catch block. This child entry has one of the tags [DW\\_TAG\\_formal\\_parameter](#) or [DW\\_TAG\\_unspecified\\_parameters](#), and will have the same form as other parameter entries.

The siblings immediately following a try block entry are its corresponding catch block entries.

## 3.9 Declarations with Reduced Scope

Any debugging information entry for a declaration (including objects, subprograms, types and modules) whose scope has an address range that is a subset of the address range for the lexical scope most closely enclosing the declared entity may have a [DW\\_AT\\_start\\_scope](#) attribute to specify that reduced range of addresses.



## Chapter 3. Program Scope Entries

1 There are two cases:

- 2 1. If the address range for the scope of the entry includes all of addresses for the  
3 containing scope except for a contiguous sequence of bytes at the beginning  
4 of the address range for the containing scope, then the address is specified  
5 using a value of class `constant`.
  - 6 a) If the address range of the containing scope is contiguous, the value of  
7 this attribute is the offset in bytes of the beginning of the address range  
8 for the scope of the object from the low PC value of the debugging  
9 information entry that defines that containing scope.
  - 10 b) If the address range of the containing scope is non-contiguous (see [2.17.3](#)  
11 [on page 53](#)) the value of this attribute is the offset in bytes of the  
12 beginning of the address range for the scope of the entity from the  
13 beginning of the first range list entry for the containing scope that is not a  
14 base address entry or an end-of-list entry.
- 15 2. Otherwise, the set of addresses for the scope of the entity is specified using a  
16 value of class `rnglistsptr`. This value indicates the beginning of a range list  
17 (see Section [2.17.3 on page 53](#)).

18 *For example, the scope of a variable may begin somewhere in the midst of a lexical `block`*  
19 *in a language that allows executable code in a block before a variable declaration, or where*  
20 *one declaration containing initialization code may change the scope of a subsequent*  
21 *declaration.*

22 Consider the following example C code:

```
float x = 99.99;  
int myfunc()  
{  
    float f = x;  
    float x = 88.99;  
    return 0;  
}
```

23 *C scoping rules require that the value of the variable `x` assigned to the variable `f` in the*  
24 *initialization sequence is the value of the global variable `x`, rather than the local `x`,*  
25 *because the scope of the local variable `x` only starts after the full declarator for the local `x`.*

26 *Due to optimization, the scope of an object may be non-contiguous and require use of a*  
27 *range list even when the containing scope is contiguous. Conversely, the scope of an*  
28 *object may not require its own range list even when the containing scope is*  
29 *non-contiguous.*



# Chapter 4

## Data Object and Object List Entries

This section presents the debugging information entries that describe individual data objects: variables, parameters and constants, and lists of those objects that may be grouped in a single declaration, such as a [common block](#).

### 4.1 Data Object Entries

Program variables, formal parameters and constants are represented by debugging information entries with the tags [DW\\_TAG\\_variable](#), [DW\\_TAG\\_formal\\_parameter](#) and [DW\\_TAG\\_constant](#), respectively.

*The tag [DW\\_TAG\\_constant](#) is used for languages that have true named constants.*

The debugging information entry for a program variable, formal parameter or constant may have the following attributes:

1. A [DW\\_AT\\_name](#) attribute, whose value is a null-terminated string containing the data object name.

If a variable entry describes an anonymous object (for example an anonymous union), the name attribute is omitted or its value consists of a single zero byte.

2. A [DW\\_AT\\_external](#) attribute, which is a [flag](#), if the name of a variable is visible outside of its enclosing compilation unit.

*The definitions of C++ static data members of structures or classes are represented by variable entries flagged as external. Both file static and local variables in C and C++ are represented by non-external variable entries.*

3. A [DW\\_AT\\_declaration](#) attribute, which is a [flag](#) that indicates whether this entry represents a non-defining declaration of an object.

## Chapter 4. Data Object and Object List

- 1 4. A [DW\\_AT\\_location](#) attribute, whose value describes the location of a variable  
2 or parameter at run-time.

3 If no location attribute is present in a variable entry representing the  
4 definition of a variable (that is, with no [DW\\_AT\\_declaration](#) attribute), or if  
5 the location attribute is present but has an empty location description (as  
6 described in Section 2.6 on page 38), the variable is assumed to exist in the  
7 source code but not in the executable program (but see number 10, below).

8 In a variable entry representing a non-defining declaration of a variable, the  
9 location specified supersedes the location specified by the defining  
10 declaration but only within the scope of the variable entry; if no location is  
11 specified, then the location specified in the defining declaration applies.

12 *This can occur, for example, for a C or C++ external variable (one that is defined and  
13 allocated in another compilation unit) and whose location varies in the current unit  
14 due to optimization.*

- 15 5. A [DW\\_AT\\_type](#) attribute describing the type of the variable, constant or  
16 formal parameter.

- 17 6. If the variable entry represents the defining declaration for a C++ static data  
18 member of a structure, class or union, the entry has a [DW\\_AT\\_specification](#)  
19 attribute, whose value is a [reference](#) to the debugging information entry  
20 representing the declaration of this data member. The referenced entry also  
21 has the tag [DW\\_TAG\\_variable](#) and will be a child of some class, structure or  
22 union type entry.

23 If the variable entry represents a non-defining declaration,  
24 [DW\\_AT\\_specification](#) may be used to reference the defining declaration of  
25 the variable. If no [DW\\_AT\\_specification](#) attribute is present, the defining  
26 declaration may be found as a global definition either in the current  
27 compilation unit or in another compilation unit with the [DW\\_AT\\_external](#)  
28 attribute.

29 Variable entries containing the [DW\\_AT\\_specification](#) attribute do not need to  
30 duplicate information provided by the declaration entry referenced by the  
31 specification attribute. In particular, such variable entries do not need to  
32 contain attributes for the name or type of the data member whose definition  
33 they represent.

- 34 7. A [DW\\_AT\\_variable\\_parameter](#) attribute, which is a [flag](#), if a formal  
35 parameter entry represents a parameter whose value in the calling function  
36 may be modified by the callee. The absence of this attribute implies that the  
37 parameter's value in the calling function cannot be modified by the callee.

## Chapter 4. Data Object and Object List

- 1 8. A **DW\_AT\_is\_optional** attribute, which is a **flag**, if a parameter entry  
2 represents an optional parameter.
- 3 9. A **DW\_AT\_default\_value** attribute for a formal parameter entry. The value of  
4 this attribute may be a constant, a reference to the debugging information  
5 entry for a variable, a reference to a debugging information entry for a  
6 DWARF procedure, or a string containing a source language fragment.
- 7 • If the attribute form is of class **constant**, that constant is interpreted as a  
8 value whose type is the same as the type of the formal parameter.  
9 *For a constant form there is no way to express the absence of a default value.*
  - 10 • If the attribute form is of class **reference**, and the referenced entry is for a  
11 variable, the default value of the parameter is the value of the referenced  
12 variable. If the reference value is 0, no default value has been specified.
  - 13 • If the attribute form is of class **string**, that string is interpreted as an  
14 expression in the source language, as defined by the compilation unit's  
15 **DW\_AT\_language\_name** and **DW\_AT\_language\_version** attributes, that  
16 is to be evaluated according to the rules defined by that source language.  
17 *The source language fragment may be different from the actual source text if the  
18 latter contains macros which have been expanded.*
- 19 10. A **DW\_AT\_const\_value** attribute for an entry describing a variable or formal  
20 parameter whose value is constant and not represented by an object in the  
21 address space of the program, or an entry describing a named constant. (Note  
22 that such an entry does not have a location attribute.) The value of this  
23 attribute may be a string or any of the constant data or data block forms, as  
24 appropriate for the representation of the variable's value. The value is the  
25 actual constant value of the variable, represented as it would be on the target  
26 architecture.  
27 *One way in which a formal parameter with a constant value and no location can arise  
28 is for a formal parameter of an inlined subprogram that corresponds to a constant  
29 actual parameter of a call that is inlined.*

11. A `DW_AT_endianity` attribute, whose value is a constant that specifies the endianness of the object. The value of this attribute specifies an ABI-defined byte ordering for the value of the object. If omitted, the default endianness of data for the given type is assumed.

The set of values and their meaning for this attribute is given in Table 4.1. These represent the default encoding formats as defined by the target architecture's ABI or processor definition. The exact definition of these formats may differ in subtle ways for different architectures.

Table 4.1: Endianness attribute values

Name	Meaning
<code>DW_END_default</code>	Default endian encoding (equivalent to the absence of a <code>DW_AT_endianity</code> attribute)
<code>DW_END_big</code>	Big-endian encoding
<code>DW_END_little</code>	Little-endian encoding

12. A `DW_AT_const_expr` attribute, constant expression attribute which is a flag, if a variable entry represents a C++ object declared with the `constexpr` specifier. This attribute indicates that the variable can be evaluated as a compile-time constant.

*In C++, a variable declared with `constexpr` is implicitly `const`. Such a variable has a `DW_AT_type` attribute whose value is a reference to a debugging information entry describing a `const` qualified type.*

13. A `DW_AT_linkage_name` attribute for a variable or constant entry as described in Section 2.22 on page 58.

## 4.2 Common Block Entries

A Fortran common block may be described by a debugging information entry with the tag `DW_TAG_common_block`.

The common block entry has a `DW_AT_name` attribute whose value is a null-terminated string containing the common block name. It may also have a `DW_AT_linkage_name` attribute as described in Section 2.22 on page 58.

A common block entry also has a `DW_AT_location` attribute whose value describes the location of the beginning of the common block.

The common block entry owns debugging information entries describing the variables contained within the common block.

1 *Fortran allows each declarer of a common block to independently define its contents;*  
2 *thus, common blocks are not types.*

### 3 **4.3 Namelist Entries**

4 *At least one language, Fortran 90, has the concept of a namelist. A namelist is an ordered*  
5 *list of the names of some set of declared objects. The namelist object itself may be used as*  
6 *a replacement for the list of names in various contexts.*

7 A namelist is represented by a debugging information entry with the tag  
8 **DW\_TAG\_namelist**. If the namelist itself has a name, the namelist entry has a  
9 **DW\_AT\_name** attribute, whose value is a null-terminated string containing the  
10 namelist's name.

11 Each name that is part of the namelist is represented by a debugging information  
12 entry with the tag **DW\_TAG\_namelist\_item**. Each such entry is a child of the  
13 namelist entry, and all of the namelist item entries for a given namelist are  
14 ordered as were the list of names they correspond to in the source program.

15 Each namelist item entry contains a **DW\_AT\_namelist\_item** attribute whose  
16 value is a **reference** to the debugging information entry representing the  
17 declaration of the item whose name appears in the namelist.

## Chapter 4. Data Object and Object List

*(empty page)*

## Chapter 4. Data Object and Object List

1 typedef

# Chapter 5

## Type Entries

This section presents the debugging information entries that describe program types: base types, modified types and user-defined types.

### 5.1 Base Type Entries

*A base type is a data type that is not defined in terms of other data types. Each programming language has a set of base types that are considered to be built into that language.*

A base type is represented by a debugging information entry with the tag `DW_TAG_base_type`.

A base type entry may have a `DW_AT_name` attribute whose value is a null-terminated string containing the name of the base type as recognized by the programming language of the compilation unit containing the base type entry.

A base type entry has a `DW_AT_encoding` attribute describing how the base type is encoded and is to be interpreted. The `DW_AT_encoding` attribute is described in Section 5.1.1 following.

A base type entry may have a `DW_AT_endianity` attribute as described in Section 4.1 on page 99. If omitted, the encoding assumes the representation that is the default for the target architecture.

A base type entry has a `DW_AT_byte_size` attribute or a `DW_AT_bit_size` attribute whose integer constant value (see Section 2.21 on page 58) is the amount of storage needed to hold a value of the type.



## Chapter 5. Type Entries

1 For example, the C type `int` on a machine that uses 32-bit integers is represented by a  
2 base type entry with a name attribute whose value is “`int`”, an encoding attribute whose  
3 value is `DW_ATE_signed` and a byte size attribute whose value is 4.

4 If the value of an object of the given type does not fully occupy the storage  
5 described by a byte size attribute, the base type entry may also have a  
6 `DW_AT_bit_size` and a `DW_AT_data_bit_offset` attribute, both of whose values  
7 are integer constant values (see Section 2.19 on page 56). The bit size attribute  
8 describes the actual size in bits used to represent values of the given type. The  
9 data bit offset attribute is the offset in bits from the beginning of the containing  
10 storage to the beginning of the value. Bits that are part of the offset are padding.  
11 If this attribute is omitted a default data bit offset of zero is assumed.

12 A `DW_TAG_base_type` entry may have additional attributes that augment  
13 certain of the base type encodings; these are described in the following section.

### 14 5.1.1 Base Type Encodings

15 A base type entry has a `DW_AT_encoding` attribute describing how the base type  
16 is encoded and is to be interpreted. The value of this attribute is an integer of  
17 class constant. The set of values and their meanings for the `DW_AT_encoding`  
18 attribute is given in Table 5.1 on the following page.

19 *In Table 5.1, encodings are shown in groups that have similar characteristics purely for*  
20 *presentation purposes. These groups are not part of this DWARF specification.*

#### 21 5.1.1.1 Simple Encodings

22 Types with simple encodings are widely supported in many programming  
23 languages and are not discussed further.

24 For a type with simple encodings, the type entry may have a `DW_AT_bias`  
25 attribute whose value is an integer constant which is added to the encoded value  
26 to determine the value of an object of the type in the source program. If the  
27 `DW_AT_bias` is encoded using `DW_FORM_data<n>`, then the bias value is  
28 treated as an unsigned integer.

## Chapter 5. Type Entries

Table 5.1: Encoding attribute values

Name	Meaning
<i>Simple encodings</i>	
DW_ATE_boolean	true or false
DW_ATE_address	machine address
DW_ATE_signed	signed binary integer
DW_ATE_signed_char	signed character
DW_ATE_unsigned	unsigned binary integer
DW_ATE_unsigned_char	unsigned character
<i>Character encodings</i>	
DW_ATE_ASCII	ISO/IEC 646:1991 character
DW_ATE_UCS	ISO/IEC 10646-1:1993 character (UCS-4)
DW_ATE_UTF	ISO/IEC 10646-1:1993 character
<i>Bit-precise integer types</i>	
DW_ATE_signed_bitint	bit-precise signed integer
DW_ATE_unsigned_bitint	bit-precise unsigned integer
<i>Scaled encodings</i>	
DW_ATE_signed_fixed	signed fixed-point scaled integer
DW_ATE_unsigned_fixed	unsigned fixed-point scaled integer
<i>Floating-point encodings</i>	
DW_ATE_float	binary floating-point number
DW_ATE_complex_float	complex binary floating-point number
DW_ATE_imaginary_float	imaginary binary floating-point number
DW_ATE_decimal_float	IEEE 754R decimal floating-point number
<i>Decimal string encodings</i>	
DW_ATE_packed_decimal	packed decimal number
DW_ATE_numeric_string	numeric string
DW_ATE_edited	edited string
<i>Complex integral encodings</i>	
DW_ATE_complex_signed	complex (signed) binary integral number
DW_ATE_imaginary_signed	imaginary (signed) binary integral number
DW_ATE_complex_unsigned	complex unsigned binary integral number
DW_ATE_imaginary_unsigned	imaginary unsigned binary integral number

1 **5.1.1.2 Character Encodings**

2 [DW\\_ATE\\_UTF](#) specifies the Unicode string encoding (see the Universal  
3 Character Set standard, ISO/IEC 10646-1:1993).

4 *For example, the C++ type `char16_t` is represented by a base type entry with a name*  
5 *attribute whose value is “`char16_t`”, an encoding attribute whose value is*  
6 *[DW\\_ATE\\_UTF](#) and a byte size attribute whose value is 2.*

7 [DW\\_ATE\\_ASCII](#) and [DW\\_ATE\\_UCS](#) specify encodings for the Fortran 2003  
8 string kinds ASCII (ISO/IEC 646:1991) and ISO 10646 (UCS-4 in ISO/IEC  
9 10646:2000).

10 **5.1.2 Bit-precise integer types**

11 Bit-precise integer types [DW\\_ATE\\_signed\\_bitint](#) and [DW\\_ATE\\_unsigned\\_bitint](#)  
12 are supported in C23<sup>1</sup>, where they are known as `_BitInt(N)` and  
13 `unsigned _BitInt(N)`, respectively.

14 **5.1.2.1 Scaled Encodings**

15 The [DW\\_ATE\\_signed\\_fixed](#) and [DW\\_ATE\\_unsigned\\_fixed](#) entries describe  
16 signed and unsigned fixed-point binary data types, respectively.

17 The fixed binary type encodings have a [DW\\_AT\\_digit\\_count](#) attribute with the  
18 same interpretation as described for the [DW\\_ATE\\_packed\\_decimal](#) and  
19 [DW\\_ATE\\_numeric\\_string](#) base type encodings (see Section 5.1.2.3 on the next  
20 page).

21 For a data type with a decimal scale factor, the fixed binary type entry has a  
22 [DW\\_AT\\_decimal\\_scale](#) attribute with the same interpretation as described for  
23 the [DW\\_ATE\\_packed\\_decimal](#) and [DW\\_ATE\\_numeric\\_string](#) base types (see  
24 Section 5.1.2.3 on the following page).

25 For a data type with a binary scale factor, the fixed binary type entry has a  
26 [DW\\_AT\\_binary\\_scale](#) attribute. The [DW\\_AT\\_binary\\_scale](#) attribute is an [integer](#)  
27 [constant](#) value that represents the exponent of the base two scale factor to be  
28 applied to an instance of the type. Zero scale puts the binary point immediately  
29 to the right of the least significant bit. Positive scale moves the binary point to the  
30 right and implies that additional zero bits on the right are not stored in an  
31 instance of the type. Negative scale moves the binary point to the left; if the  
32 absolute value of the scale is larger than the number of bits, this implies  
33 additional zero bits on the left are not stored in an instance of the type.

---

<sup>1</sup>C23 is an informal name for what will likely become ISO/IEC 9899:2024.

## Chapter 5. Type Entries

1 For a data type with a rational scale factor, one or both of the following attributes  
2 may be used:

- 3 • **DW\_AT\_scale\_multiplier**. This attribute is an integer constant value that  
4 represents a multiplicative scale factor to be applied to an instance of the  
5 type.
- 6 • **DW\_AT\_scale\_divisor**. This attribute is an integer constant value that  
7 represents the reciprocal of a multiplicative scale factor to be applied to an  
8 instance of the type.

9 If both attributes are present, both are applied, with the result being equivalent to  
10 a rational scale factor  $x/y$ , where  $x$  is the value of **DW\_AT\_scale\_multiplier** and  $y$   
11 is the value of **DW\_AT\_scale\_divisor**.

12 For a data type with a non-rational scale factor, the fixed binary type entry has a  
13 **DW\_AT\_small** attribute which references a **DW\_TAG\_constant** entry. The scale  
14 factor value is interpreted in accordance with the value defined by the  
15 **DW\_TAG\_constant** entry. The value represented is the product of the integer  
16 value in memory and the associated constant entry for the type.

17 *The **DW\_AT\_small** attribute is defined with the Ada `small` attribute in mind.*

18 If a type entry has attributes that describe more than one kind of scale factor, the  
19 resulting scale factor for the type is the product of the individual scale factors.

### 20 5.1.2.2 Floating-Point Encodings

21 Types with binary floating-point encodings (**DW\_ATE\_float**,  
22 **DW\_ATE\_complex\_float** and **DW\_ATE\_imaginary\_float**) are supported in many  
23 programming languages and are not discussed further.

24 **DW\_ATE\_decimal\_float** specifies floating-point representations that have a  
25 power-of-ten exponent, such as specified in IEEE 754R.

### 26 5.1.2.3 Decimal String Encodings

27 The **DW\_ATE\_packed\_decimal** and **DW\_ATE\_numeric\_string** base type  
28 encodings represent packed and unpacked decimal string numeric data types,  
29 respectively, either of which may be either signed or unsigned. These base types  
30 are used in combination with **DW\_AT\_decimal\_sign**, **DW\_AT\_digit\_count** and  
31 **DW\_AT\_decimal\_scale** attributes.

## Chapter 5. Type Entries

1 A `DW_AT_decimal_sign` attribute is an `integer constant` that conveys the  
2 representation of the sign of the decimal type (see Table 5.2). Its `integer constant`  
3 value is interpreted to mean that the type has a leading overpunch, trailing  
4 overpunch, leading separate or trailing separate sign representation or,  
5 alternatively, no sign at all.

Table 5.2: Decimal sign attribute values

Name	Meaning
<code>DW_DS_unsigned</code>	Unsigned
<code>DW_DS_leading_overpunch</code>	Sign is encoded in the most significant digit in a target-dependent manner
<code>DW_DS_trailing_overpunch</code>	Sign is encoded in the least significant digit in a target-dependent manner
<code>DW_DS_leading_separate</code>	Decimal type: Sign is a "+" or "-" character to the left of the most significant digit.
<code>DW_DS_trailing_separate</code>	Decimal type: Sign is a "+" or "-" character to the right of the least significant digit. Packed decimal type: Least significant nibble contains a target-dependent value indicating positive or negative.

6 The `DW_AT_decimal_scale` attribute is an integer constant value that represents  
7 the exponent of the base ten scale factor to be applied to an instance of the type.  
8 A scale of zero puts the decimal point immediately to the right of the least  
9 significant digit. Positive scale moves the decimal point to the right and implies  
10 that additional zero digits on the right are not stored in an instance of the type.  
11 Negative scale moves the decimal point to the left; if the absolute value of the  
12 scale is larger than the digit count, this implies additional zero digits on the left  
13 are not stored in an instance of the type.

14 The `DW_AT_digit_count` attribute is an `integer constant` value that represents the  
15 number of digits in an instance of the type.

16 The `DW_ATE_edited` base type is used to represent an edited numeric or  
17 alphanumeric data type. It is used in combination with a `DW_AT_picture_string`  
18 attribute whose value is a null-terminated string containing the target-dependent  
19 picture string associated with the type.

1 If the edited base type entry describes an edited numeric data type, the edited  
2 type entry has a `DW_AT_digit_count` and a `DW_AT_decimal_scale` attribute.  
3 These attributes have the same interpretation as described for the  
4 `DW_ATE_packed_decimal` and `DW_ATE_numeric_string` base types. If the  
5 edited type entry describes an edited alphanumeric data type, the edited type  
6 entry does not have these attributes.

7 *The presence or absence of the `DW_AT_digit_count` and `DW_AT_decimal_scale`  
8 attributes allows a debugger to easily distinguish edited numeric from edited  
9 alphanumeric, although in principle the digit count and scale are derivable by  
10 interpreting the picture string.*

#### 11 5.1.2.4 Complex Integral Encodings

12 Complex types with binary integral encodings (`DW_ATE_complex_signed`,  
13 `DW_ATE_imaginary_signed`, `DW_ATE_complex_unsigned` and  
14 `DW_ATE_imaginary_unsigned`) are supported in some programming languages  
15 (for example, GNU C and Rust) and are not discussed further."

## 16 5.2 Unspecified Type Entries

17 Some languages have constructs in which a type may be left unspecified or the  
18 absence of a type may be explicitly indicated.

19 An unspecified (implicit, unknown, ambiguous or nonexistent) type is  
20 represented by a debugging information entry with the tag  
21 `DW_TAG_unspecified_type`. If a name has been given to the type, then the  
22 corresponding unspecified type entry has a `DW_AT_name` attribute whose value  
23 is a null-terminated string containing the name.

24 *The interpretation of this debugging information entry is intentionally left flexible to  
25 allow it to be interpreted appropriately in different languages. For example, in C and  
26 C++ the language implementation can provide an unspecified type entry with the name  
27 "void" which can be referenced by the type attribute of pointer types and typedef  
28 declarations for 'void' (see Sections 5.3 on the following page and 5.4 on page 115,  
29 respectively). As another example, in Ada such an unspecified type entry can be referred  
30 to by the type attribute of an access type where the denoted type is incomplete (the name  
31 is declared as a type but the definition is deferred to a separate compilation unit).*

32 *C++ permits using the `auto` return type specifier for the return type of a member  
33 function declaration. The actual return type is deduced based on the definition of the  
34 function, so it may not be known when the function is declared. The language  
35 implementation can provide an unspecified type entry with the name `auto` which can be*

1 *referenced by the return type attribute of a function declaration entry. When the function*  
 2 *is later defined, the [DW\\_TAG\\_subprogram](#) entry for the definition includes a reference to*  
 3 *the actual return type.*

## 4 **5.3 Type Modifier Entries**

5 A base or user-defined type may be modified in different ways in different  
 6 languages. A type modifier is represented in DWARF by a debugging  
 7 information entry with one of the tags given in Table 5.3.

Table 5.3: Type modifier tags

Name	Meaning
<a href="#">DW_TAG_atomic_type</a>	atomic qualified type (for example, in C)
<a href="#">DW_TAG_const_type</a>	const qualified type (for example in C, C++)
<a href="#">DW_TAG_immutable_type</a>	immutable type (for example, in D)
<a href="#">DW_TAG_packed_type</a>	packed type (for example in Ada, Pascal)
<a href="#">DW_TAG_pointer_type</a>	pointer to an object of the type being modified
<a href="#">DW_TAG_reference_type</a>	reference to (lvalue of) an object of the type being modified
<a href="#">DW_TAG_restrict_type</a>	restrict qualified type
<a href="#">DW_TAG_rvalue_reference_type</a>	rvalue reference to an object of the type being modified (for example, in C++)
<a href="#">DW_TAG_shared_type</a>	shared qualified type (for example, in UPC)
<a href="#">DW_TAG_volatile_type</a>	volatile qualified type (for example, in C, C++)

8 If a name has been given to the modified type in the source program, then the  
 9 corresponding modified type entry has a [DW\\_AT\\_name](#) attribute whose value is  
 10 a null-terminated string containing the name of the modified type.

11 Each of the type modifier entries has a [DW\\_AT\\_type](#) attribute, whose value is a  
 12 [reference](#) to a debugging information entry describing a base type, a user-defined  
 13 type or another type modifier.



## Chapter 5. Type Entries

As examples of how type modifiers are ordered, consider the following C declarations:

```
const unsigned char * volatile p;
```

This represents a volatile pointer to a constant character. It is encoded in DWARF as

```
DW_TAG_variable(p) -->
  DW_TAG_volatile_type -->
    DW_TAG_pointer_type -->
      DW_TAG_const_type -->
        DW_TAG_base_type(unsigned char)
```

On the other hand

```
volatile unsigned char * const restrict p;
```

represents a restricted constant pointer to a volatile character. This is encoded as

```
DW_TAG_variable(p) -->
  DW_TAG_restrict_type -->
    DW_TAG_const_type -->
      DW_TAG_pointer_type -->
        DW_TAG_volatile_type -->
          DW_TAG_base_type(unsigned char)
```

Figure 5.1: Type modifier examples

1 A modified type entry describing a pointer or reference type (using  
2 [DW\\_TAG\\_pointer\\_type](#), [DW\\_TAG\\_reference\\_type](#) or  
3 [DW\\_TAG\\_rvalue\\_reference\\_type](#)) may have a [DW\\_AT\\_address\\_class](#) attribute to  
4 describe how objects having the given pointer or reference type are dereferenced.

5 A modified type entry describing a UPC shared qualified type (using  
6 [DW\\_TAG\\_shared\\_type](#)) may have a [DW\\_AT\\_count](#) attribute whose value is a  
7 constant expressing the (explicit or implied) blocksize specified for the type in the  
8 source. If no count attribute is present, then the “infinite” blocksize is assumed.

9 When multiple type modifiers are chained together to modify a base or  
10 user-defined type, the tree ordering reflects the semantics of the applicable  
11 language rather than the textual order in the source presentation.

12 Examples of modified types are shown in Figure 5.1.



## 5.4 Typedef Entries

A named type that is defined in terms of another type definition is represented by a debugging information entry with the tag `DW_TAG_typedef`. The typedef entry has a `DW_AT_name` attribute whose value is a null-terminated string containing the name of the typedef.

The typedef entry may also contain a `DW_AT_type` attribute whose value is a [reference](#) to the type named by the typedef. If the debugging information entry for a typedef represents a declaration of the type that is not also a definition, it does not contain a type attribute.

*Depending on the language, a named type that is defined in terms of another type may be called a type alias, a subtype, a constrained type and other terms. A type name declared with no defining details may be termed an incomplete, forward or hidden type. While the DWARF `DW_TAG_typedef` entry was originally inspired by the like named construct in C and C++, it is broadly suitable for similar constructs (by whatever source syntax) in other languages.*

## 5.5 Array Type Entries

*Many languages share the concept of an “array,” which is a table of components of identical type. Furthermore, many architectures contain vector types which mirror the language concept of a short single dimension array but have different encoding, a different calling convention and different arithmetic and logical operational semantics than the source language arrays. Likewise, a few architectures are starting to add matrix register types with similar variations in encoding and semantics from normal source language array types.*

An array type is represented by a debugging information entry with the tag `DW_TAG_array_type`. If a name has been given to the array type in the source program, then the corresponding array type entry has a `DW_AT_name` attribute whose value is a null-terminated string containing the array type name.

The array type may have a `DW_AT_tensor` attribute, which is a flag. If present, this attribute indicates that the entry describes a vector or matrix type. The array dimensions (see below) describe the vector width, and when applicable the number of rows.

The array type entry describing a multidimensional array may have a `DW_AT_ordering` attribute whose [integer constant](#) value is interpreted to mean either row-major or column-major ordering of array elements. The set of values and their meanings for the ordering attribute are listed in Table 5.4 following. If

## Chapter 5. Type Entries

1 no ordering attribute is present, the default ordering for the source language  
2 (which is indicated by the [DW\\_AT\\_language\\_name](#) attribute of the enclosing  
3 compilation unit entry) is assumed.

Table 5.4: Array ordering

---

[DW\\_ORD\\_col\\_major](#)  
[DW\\_ORD\\_row\\_major](#)

---

4 An array type entry has a [DW\\_AT\\_type](#) attribute describing the type of each  
5 element of the array. If [DW\\_AT\\_tensor](#) is present, the element type must be a  
6 base type (see Section 5.1 on page 106).

7 If the amount of storage allocated to hold each element of an object of the given  
8 array type is different from the amount of storage that is normally allocated to  
9 hold an individual object of the indicated element type, then the array type entry  
10 has either a [DW\\_AT\\_byte\\_stride](#) or a [DW\\_AT\\_bit\\_stride](#) attribute, whose value  
11 (see Section 2.19 on page 56) is the size of each element of the array.

12 The array type entry may have either a [DW\\_AT\\_byte\\_size](#) or a [DW\\_AT\\_bit\\_size](#)  
13 attribute (see Section 2.21 on page 58), whose value is the amount of storage  
14 needed to hold an instance of the array type.

15 *If the size of the array can be determined statically at compile time, this value can usually  
16 be computed by multiplying the number of array elements by the size of each element.*

17 Each array dimension is described by a debugging information entry with either  
18 the tag [DW\\_TAG\\_subrange\\_type](#) or the tag [DW\\_TAG\\_enumeration\\_type](#). These  
19 entries are children of the array type entry and are ordered to reflect the  
20 appearance of the dimensions in the source program (that is, leftmost dimension  
21 first, next to leftmost second, and so on).

22 *In languages that have no concept of a “multidimensional array” (for example, C), an  
23 array of arrays may be represented by a debugging information entry for a  
24 multidimensional array.*

1 Alternatively, for an array with dynamic rank the array dimensions are described  
2 by a debugging information entry with the tag [DW\\_TAG\\_generic\\_subrange](#).  
3 This entry has the same attributes as a [DW\\_TAG\\_subrange\\_type](#) entry; however,  
4 there is just one [DW\\_TAG\\_generic\\_subrange](#) entry and it describes all of the  
5 dimensions of the array. If [DW\\_TAG\\_generic\\_subrange](#) is used, the number of  
6 dimensions must be specified using a [DW\\_AT\\_rank](#) attribute. See also Section  
7 [5.18.3 on page 139](#).

8 Other attributes especially applicable to arrays are [DW\\_AT\\_allocated](#),  
9 [DW\\_AT\\_associated](#) and [DW\\_AT\\_data\\_location](#), which are described in Section  
10 [5.18 on page 137](#). For relevant examples, see also Appendix [D.2.1 on page 304](#).

## 11 5.6 Coarray Type Entries

12 *In Fortran, a “coarray” is an array whose elements are located in different processes*  
13 *rather than in the memory of one process. The individual elements of a coarray can be*  
14 *scalars or arrays. Similar to arrays, coarrays have “codimensions” that are indexed using*  
15 *a “coindex” or multiple “coindices”.*

16 A coarray type is represented by a debugging information entry with the tag  
17 [DW\\_TAG\\_coarray\\_type](#). If a name has been given to the coarray type in the  
18 source, then the corresponding coarray type entry has a [DW\\_AT\\_name](#) attribute  
19 whose value is a null-terminated string containing the array type name.

20 A coarray entry has one or more [DW\\_TAG\\_subrange\\_type](#) child entries, one for  
21 each codimension. It also has a [DW\\_AT\\_type](#) attribute describing the type of  
22 each element of the coarray.

23 *In a coarray application, the run-time number of processes in the application is part of the*  
24 *coindex calculation. It is represented in the Fortran source by a coindex which is declared*  
25 *with a “\*” as the upper bound. To express this concept in DWARF, the*  
26 *[DW\\_TAG\\_subrange\\_type](#) child entry for that index has only a lower bound and no*  
27 *upper bound.*

28 *How coarray elements are located and how coindices are converted to process*  
29 *specifications is implementation-defined.*

## 5.7 Structure, Union, Class and Interface Type Entries

*The languages C, C++, and Pascal, among others, allow the programmer to define types that are collections of related components. In C and C++, these collections are called "structures." In Pascal, they are called "records." The components may be of different types. The components are called "members" in C and C++, and "fields" in Pascal.*

*The components of these collections each exist in their own space in computer memory. The components of a C or C++ "union" all coexist in the same memory.*

*Pascal and other languages have a "discriminated union," also called a "variant record." Here, selection of a number of alternative substructures ("variants") is based on the value of a component that is not part of any of those substructures (the "discriminant").*

*C++ and Java have the notion of "class," which is in some ways similar to a structure. A class may have "member functions" which are subroutines that are within the scope of a class or structure.*

*The C++ notion of structure is more general than in C, being equivalent to a class with minor differences. Accordingly, in the following discussion, statements about C++ classes may be understood to apply to C++ structures as well.*

### 5.7.1 Structure, Union and Class Type Entries

Structure, union, and class types are represented by debugging information entries with the tags `DW_TAG_structure_type`, `DW_TAG_union_type`, and `DW_TAG_class_type`, respectively. If a name has been given to the structure, union, or class in the source program, then the corresponding structure type, union type, or class type entry has a `DW_AT_name` attribute whose value is a null-terminated string containing the type name.

The members of a structure, union, or class are represented by debugging information entries that are owned by the corresponding structure type, union type, or class type entry and appear in the same order as the corresponding declarations in the source program.

A structure, union, or class type may have a `DW_AT_export_symbols` attribute which indicates that all member names defined within the structure, union, or class may be referenced as if they were defined within the containing structure, union, or class.

*This may be used to describe anonymous structures, unions and classes in C or C++.*

## Chapter 5. Type Entries

1 A structure type, union type or class type entry may have either a  
2 [DW\\_AT\\_byte\\_size](#) or a [DW\\_AT\\_bit\\_size](#) attribute (see Section 2.21 on page 58),  
3 whose value is the amount of storage needed to hold an instance of the structure,  
4 union or class type, including any padding.

5 An incomplete structure, union or class type is represented by a structure, union  
6 or class entry that does not have a byte size attribute and that has a  
7 [DW\\_AT\\_declaration](#) attribute.

8 If the complete declaration of a type has been placed in a separate type unit (see  
9 Section 3.1.4 on page 71), an incomplete declaration of that type in the  
10 compilation unit may provide the unique 8-byte signature of the type using a  
11 [DW\\_AT\\_signature](#) attribute.

12 If a structure, union or class entry represents the definition of a structure, union  
13 or class member corresponding to a prior incomplete structure, union or class,  
14 the entry may have a [DW\\_AT\\_specification](#) attribute whose value is a [reference](#)  
15 to the debugging information entry representing that incomplete declaration.

16 Structure, union and class entries containing the [DW\\_AT\\_specification](#) attribute  
17 do not need to duplicate information provided by the declaration entry  
18 referenced by the specification attribute. In particular, such entries do not need to  
19 contain an attribute for the name of the structure, union or class they represent if  
20 such information is already provided in the declaration.

21 *For C and C++, data member declarations occurring within the declaration of a*  
22 *structure, union or class type are considered to be “definitions” of those members, with*  
23 *the exception of “static” data members, whose definitions appear outside of the*  
24 *declaration of the enclosing structure, union or class type. Function member declarations*  
25 *appearing within a structure, union or class type declaration are definitions only if the*  
26 *body of the function also appears within the type declaration.*

27 If the definition for a given member of the structure, union or class does not  
28 appear within the body of the declaration, that member also has a debugging  
29 information entry describing its definition. That latter entry has a  
30 [DW\\_AT\\_specification](#) attribute referencing the debugging information entry  
31 owned by the body of the structure, union or class entry and representing a  
32 non-defining declaration of the data, function or type member. The referenced  
33 entry will not have information about the location of that member (low and high  
34 PC attributes for function members, location descriptions for data members) and  
35 will have a [DW\\_AT\\_declaration](#) attribute.

## Chapter 5. Type Entries

1 Consider a nested class whose definition occurs outside of the containing class definition,  
2 as in:

```
struct A {  
    struct B;  
};  
struct A::B { ... };
```

3 The two different structs can be described in different compilation units to facilitate  
4 DWARF space compression (see [Appendix E.1 on page 386](#)).

5 A structure type, union type or class type entry may have a  
6 **DW\_AT\_calling\_convention** attribute, whose value indicates whether a value of  
7 the type is passed by reference or passed by value. The set of calling convention  
8 codes for use with types is given in [Table 5.5](#) following.

Table 5.5: Calling convention codes for types

---

[DW\\_CC\\_normal](#)  
[DW\\_CC\\_pass\\_by\\_value](#)  
[DW\\_CC\\_pass\\_by\\_reference](#)

---

9 If this attribute is not present, or its value is [DW\\_CC\\_normal](#), the convention to  
10 be used for an object of the given type is assumed to be unspecified.

11 Note that [DW\\_CC\\_normal](#) is also used as a calling convention code for certain  
12 subprograms (see [Table 3.4 on page 79](#)).

13 If unspecified, a consumer may be able to deduce the calling convention based on  
14 knowledge of the type and the ABI.

### 15 5.7.2 Interface Type Entries

16 The Java language defines “interface” types. An interface in Java is similar to a C++ or  
17 Java class with only abstract methods and constant data members.

18 Interface types are represented by debugging information entries with the tag  
19 **DW\_TAG\_interface\_type**.

20 An interface type entry has a [DW\\_AT\\_name](#) attribute, whose value is a  
21 null-terminated string containing the type name.

22 The members of an interface are represented by debugging information entries  
23 that are owned by the interface type entry and that appear in the same order as  
24 the corresponding declarations in the source program.

### 5.7.3 Derived or Extended Structures, Classes and Interfaces

*In C++, a class (or struct) may be “derived from” or be a “subclass of” another class. In Java, an interface may “extend” one or more other interfaces, and a class may “extend” another class and/or “implement” one or more interfaces. All of these relationships may be described using the following. Note that in Java, the distinction between extends and implements is implied by the entities at the two ends of the relationship.*

A class type or interface type entry that describes a derived, extended or implementing class or interface owns debugging information entries describing each of the classes or interfaces it is derived from, extending or implementing, respectively, ordered as they were in the source program. Each such entry has the tag **DW\_TAG\_inheritance**.

An inheritance entry has a **DW\_AT\_type** attribute whose value is a reference to the debugging information entry describing the class or interface from which the parent class or structure of the inheritance entry is derived, extended or implementing.

An inheritance entry for a class that derives from or extends another class or struct also has a **DW\_AT\_data\_member\_location** attribute, whose value describes the location of the beginning of the inherited type relative to the beginning address of the instance of the derived class. If that value is a constant, it is the offset in bytes from the beginning of the class to the beginning of the instance of the inherited type. Otherwise, the value must be a location description. In this latter case, the beginning address of the instance of the derived class is pushed on the expression stack before the location description is evaluated and the result of the evaluation is the location of the instance of the inherited type.

*The interpretation of the value of this attribute for inherited types is the same as the interpretation for data members (see Section 5.7.6 following).*

An inheritance entry may have a **DW\_AT\_accessibility** attribute. If no accessibility attribute is present, private access is assumed for an entry of a class and public access is assumed for an entry of a struct, union or interface.

If the class referenced by the inheritance entry serves as a C++ virtual base class, the inheritance entry has a **DW\_AT\_virtuality** attribute.

*For a C++ virtual base, the data member location attribute will usually consist of a non-trivial location description.*



## 5.7.4 Access Declarations

*In C++, a derived class may contain access declarations that change the accessibility of individual class members from the overall accessibility specified by the inheritance declaration. A single access declaration may refer to a set of overloaded names.*

If a derived class or structure contains access declarations, each such declaration may be represented by a debugging information entry with the tag `DW_TAG_access_declaration`. Each such entry is a child of the class or structure type entry.

An access declaration entry has a `DW_AT_name` attribute, whose value is a null-terminated string representing the name used in the declaration, including any class or structure qualifiers.

An access declaration entry also has a `DW_AT_accessibility` attribute describing the declared accessibility of the named entities.

## 5.7.5 Friends

Each friend declared by a structure, union or class type may be represented by a debugging information entry that is a child of the structure, union or class type entry; the friend entry has the tag `DW_TAG_friend`.

A friend entry has a `DW_AT_friend` attribute, whose value is a reference to the debugging information entry describing the declaration of the friend.

## 5.7.6 Data Member Entries

A data member (as opposed to a member function) is represented by a debugging information entry with the tag `DW_TAG_member`. The member entry for a named member has a `DW_AT_name` attribute whose value is a null-terminated string containing the member name. If the member entry describes an anonymous union, the name attribute is omitted or the value of the attribute consists of a single zero byte.

The data member entry has a `DW_AT_type` attribute to denote the type of that member.

A data member entry may have a `DW_AT_accessibility` attribute. If no accessibility attribute is present, private access is assumed for an member of a class and public access is assumed for an member of a structure, union, or interface.



## Chapter 5. Type Entries

1 A data member entry may have a `DW_AT_mutable` attribute, which is a `flag`.  
2 This attribute indicates whether the data member was declared with the mutable  
3 storage class specifier.

4 The beginning of a data member is described relative to the beginning of the  
5 object in which it is immediately contained. In general, the beginning is  
6 characterized by both an address and a bit offset within the byte at that address.  
7 When the storage for an entity includes all of the bits in the beginning byte, the  
8 beginning bit offset is defined to be zero.

9 The member entry corresponding to a data member that is defined in a structure,  
10 union or class may have either a `DW_AT_data_member_location` attribute or a  
11 `DW_AT_data_bit_offset` attribute. If the beginning of the data member is the  
12 same as the beginning of the containing entity then neither attribute is required.

13 For a `DW_AT_data_member_location` attribute there are two cases:

- 14 1. If the value is an `integer constant`, it is the offset in bytes from the beginning  
15 of the containing entity. If the beginning of the containing entity has a  
16 non-zero bit offset then the beginning of the member entry has that same bit  
17 offset as well.
- 18 2. Otherwise, the value must be a location description. In this case, the  
19 beginning of the containing entity must be byte aligned. The beginning  
20 address is pushed on the DWARF stack before the location description is  
21 evaluated; the result of the evaluation is the base address of the member  
22 entry.

23 *The push on the DWARF expression stack of the base address of the containing*  
24 *construct is equivalent to execution of the `DW_OP_push_object_address` operation*  
25 *(see Section 2.5.1.3 on page 29); `DW_OP_push_object_address` therefore is not*  
26 *needed at the beginning of a location description for a data member. The result of the*  
27 *evaluation is a location—either an address or the name of a register, not an offset to*  
28 *the member.*

29 *A `DW_AT_data_member_location` attribute that has the form of a location*  
30 *description is not valid for a data member contained in an entity that is not byte*  
31 *aligned because DWARF operations do not allow for manipulating or computing bit*  
32 *offsets.*

33 For a `DW_AT_data_bit_offset` attribute, the value is an `integer constant` (see  
34 Section 2.19 on page 56) that specifies the number of bits from the beginning of  
35 the containing entity to the beginning of the data member. This value must be  
36 greater than or equal to zero, but is not limited to less than the number of bits per  
37 byte.

1 If the size of a data member is not the same as the size of the type given for the  
2 data member, the data member has either a [DW\\_AT\\_byte\\_size](#) or a  
3 [DW\\_AT\\_bit\\_size](#) attribute whose [integer constant](#) value (see Section 2.19 on  
4 [page 56](#)) is the amount of storage needed to hold the value of the data member.

5 *For showing nested and packed records and arrays, see Appendix D.2.7 on page 321 and*  
6 *D.2.8 on page 323.*

### 7 **5.7.7 Class Variable Entries**

8 A class variable (“static data member” in C++) is a variable shared by all  
9 instances of a class. It is represented by a debugging information entry with the  
10 tag [DW\\_TAG\\_variable](#).

11 The class variable entry may contain the same attributes and follows the same  
12 rules as non-member global variable entries (see Section 4.1 on page 99).

13 A class variable entry may have a [DW\\_AT\\_accessibility](#) attribute. If no  
14 accessibility attribute is present, private access is assumed for an entry of a class  
15 and public access is assumed for an entry of a structure, union or interface.

### 16 **5.7.8 Member Function Entries**

17 A member function is represented by a debugging information entry with the tag  
18 [DW\\_TAG\\_subprogram](#). The member function entry may contain the same  
19 attributes and follows the same rules as non-member global subroutine entries  
20 (see Section 3.3 on page 78).

21 *In particular, if the member function entry is an instantiation of a member function*  
22 *template, it follows the same rules as function template instantiations (see Section 3.3.7*  
23 *on page 84).*

24 A member function entry may have a [DW\\_AT\\_accessibility](#) attribute. If no  
25 accessibility attribute is present, private access is assumed for an entry of a class  
26 and public access is assumed for an entry of a structure, union or interface.

27 If the member function entry describes a virtual function, then that entry has a  
28 [DW\\_AT\\_virtuality](#) attribute.

29 If the member function entry describes an explicit member function, then that  
30 entry has a [DW\\_AT\\_explicit](#) attribute.

## Chapter 5. Type Entries

1 An entry for a virtual function also has a `DW_AT_vtable_elem_location` attribute  
2 whose value contains a location description yielding the address of the slot for  
3 the function within the virtual function table for the enclosing class. The address  
4 of an object of the enclosing type is pushed onto the expression stack before the  
5 location description is evaluated.

6 If the member function entry describes a non-static member function, then that  
7 entry has a `DW_AT_object_pointer` attribute whose value is a [reference](#) to the  
8 formal parameter entry that corresponds to the object for which the function is  
9 called. The name attribute of that formal parameter is defined by the current  
10 language (for example, `this` for C++ or `self` for Objective C and some other  
11 languages). That parameter also has a `DW_AT_artificial` attribute whose value is  
12 true.

13 Conversely, if the member function entry describes a static member function, the  
14 entry does not have a `DW_AT_object_pointer` attribute.

15 *In C++, non-static member functions can have const-volatile qualifiers, which affect the*  
16 *type of the first formal parameter (the “this”-pointer).*

17 If the member function entry describes a non-static member function that has a  
18 const-volatile qualification, then the entry describes a non-static member  
19 function whose object formal parameter has a type that has an equivalent  
20 const-volatile qualification.

21 *Beginning in C++11, non-static member functions can also have one of the ref-qualifiers,*  
22 *& and &&. These do not change the type of the “this”-pointer, but they do affect the*  
23 *types of object values on which the function can be invoked.*

24 The member function entry may have an `DW_AT_reference` attribute to indicate  
25 a non-static member function that can only be called on lvalue objects, or the  
26 `DW_AT_rvalue_reference` attribute to indicate that it can only be called on  
27 prvalues and xvalues.

28 *The lvalue, prvalue and xvalue concepts are defined in the C++11 and later standards.*

29 If a subroutine entry represents the defining declaration of a member function  
30 and that definition appears outside of the body of the enclosing class declaration,  
31 the subroutine entry has a `DW_AT_specification` attribute, whose value is a  
32 reference to the debugging information entry representing the declaration of this  
33 function member. The referenced entry will be a child of some class (or structure)  
34 type entry.

## Chapter 5. Type Entries

1 Subroutine entries containing the `DW_AT_specification` attribute do not need to  
2 duplicate information provided by the declaration entry referenced by the  
3 specification attribute. In particular, such entries do not need to contain a name  
4 attribute giving the name of the function member whose definition they  
5 represent. Similarly, such entries do not need to contain a return type attribute,  
6 unless the return type on the declaration was unspecified (for example, the  
7 declaration used the C++ `auto` return type specifier).

8 *In C++, a member function may be declared as deleted. This prevents the compiler from  
9 generating a default implementation of a special member function such as a constructor  
10 or destructor, and can affect overload resolution when used on other member functions.*

11 If the member function entry has been declared as deleted, then that entry has a  
12 `DW_AT_deleted` attribute.

13 *In C++, a special member function may be declared as defaulted, which explicitly declares  
14 a default compiler-generated implementation of the function. The declaration may have  
15 different effects on the calling convention used for objects of its class, depending on  
16 whether the default declaration is made inside or outside the class.*

17 If the member function has been declared as defaulted, then the entry has a  
18 `DW_AT_defaulted` attribute whose integer constant value indicates whether, and  
19 if so, how, that member is defaulted. The possible values and their meanings are  
20 shown in Table 5.6 following.

Table 5.6: Defaulted attribute names

Defaulted attribute name	Meaning
<code>DW_DEFAULTED_no</code>	Not declared default
<code>DW_DEFAULTED_in_class</code>	Defaulted within the class
<code>DW_DEFAULTED_out_of_class</code>	Defaulted outside of the class

21 *An artificial member function (that is, a compiler-generated copy that does not appear in  
22 the source) does not have a `DW_AT_defaulted` attribute.*

### 23 5.7.9 Class Template Instantiations

24 *In C++ a class template is a generic definition of a class type that may be instantiated  
25 when an instance of the class is declared or defined. The generic description of the class  
26 may include parameterized types, parameterized compile-time constant values, and/or  
27 parameterized run-time constant addresses. DWARF does not represent the generic  
28 template definition, but does represent each instantiation.*

## Chapter 5. Type Entries

1 A class template instantiation is represented by a debugging information entry  
2 with the tag `DW_TAG_class_type`, `DW_TAG_structure_type` or  
3 `DW_TAG_union_type`. With the following exceptions, such an entry will contain  
4 the same attributes and have the same types of child entries as would an entry  
5 for a class type defined explicitly using the instantiation types and values. The  
6 exceptions are:

- 7 1. Template parameters are described and referenced as specified in Section 2.23  
8 on page 58.
- 9 2. If the compiler has generated a special compilation unit to hold the template  
10 instantiation and that special compilation unit has a different name from the  
11 compilation unit containing the template definition, the name attribute for  
12 the debugging information entry representing the special compilation unit is  
13 empty or omitted.
- 14 3. If the class type entry representing the template instantiation or any of its  
15 child entries contains declaration coordinate attributes, those attributes refer  
16 to the source for the template definition, not to any source generated  
17 artificially by the compiler.

### 18 5.7.10 Variant Entries

19 A variant part of a structure is represented by a debugging information entry  
20 with the tag `DW_TAG_variant_part` and is owned by the corresponding  
21 structure type entry.

22 If the variant part has a discriminant, the discriminant is represented by a  
23 separate debugging information entry. This entry has the form of a structure data  
24 member entry. The variant part entry will have a `DW_AT_discr` attribute whose  
25 value is a [reference](#) to the member entry for the discriminant.

26 If the variant part does not have a discriminant (tag field), the variant part entry  
27 may have a `DW_AT_type` attribute to represent the tag type.

28 *A reference to a type supports the Pascal notion of a tagless variant part where the*  
29 *omitted tag nonetheless is given a type whose values are used in later parts of the variant*  
30 *syntax.*

## Chapter 5. Type Entries

1 Each variant of a particular variant part is represented by a debugging  
2 information entry with the tag `DW_TAG_variant` and is a child of the variant  
3 part entry. The value that selects a given variant may be represented in one of  
4 three ways. The variant entry may have a `DW_AT_discr_value` attribute whose  
5 value represents the discriminant value selecting this variant. The value of this  
6 attribute is encoded as an LEB128 number. The number is signed if the tag type  
7 for the variant part containing this variant is a signed type. The number is  
8 unsigned if the tag type is an unsigned type.

9 Alternatively, the variant entry may contain a `DW_AT_discr_list` attribute, whose  
10 value represents a list of discriminant values. This list is represented by any of  
11 the `block` forms and may contain a mixture of discriminant values and  
12 discriminant ranges. Each item on the list is prefixed with a discriminant value  
13 descriptor that determines whether the list item represents a single label or a  
14 label range. A single case label is represented as an LEB128 number as defined  
15 above for the `DW_AT_discr_value` attribute. A label range is represented by two  
16 LEB128 numbers, the low value of the range followed by the high value. Both  
17 values follow the rules for signedness just described. The discriminant value  
18 descriptor is an integer constant that may have one of the values given in Table  
19 5.7.

Table 5.7: Discriminant descriptor values

---

`DW_DSC_label`  
`DW_DSC_range`

---

20 If a variant entry has neither a `DW_AT_discr_value` attribute nor a  
21 `DW_AT_discr_list` attribute, or if it has a `DW_AT_discr_list` attribute with 0 size,  
22 the variant is a default variant.

23 The components selected by a particular variant are represented by debugging  
24 information entries owned by the corresponding variant entry and appear in the  
25 same order as the corresponding declarations in the source program.

26 *For examples using variant entries in several languages, see Section D.2.10 on page 326.*

## 5.8 Condition Entries

COBOL has the notion of a “level-88 condition” that associates a data item, called the conditional variable, with a set of one or more constant values and/or value ranges. Semantically, the condition is ‘true’ if the conditional variable’s value matches any of the described constants, and the condition is ‘false’ otherwise.

The **DW\_TAG\_condition** debugging information entry describes a logical condition that tests whether a given data item’s value matches one of a set of constant values. If a name has been given to the condition, the condition entry has a **DW\_AT\_name** attribute whose value is a null-terminated string giving the condition name.

The condition entry’s parent entry describes the conditional variable; normally this will be a **DW\_TAG\_variable**, **DW\_TAG\_member** or **DW\_TAG\_formal\_parameter** entry. If the parent entry has an array type, the condition can test any individual element, but not the array as a whole. The condition entry implicitly specifies a “comparison type” that is the type of an array element if the parent has an array type; otherwise it is the type of the parent entry.

The condition entry owns **DW\_TAG\_constant** and/or **DW\_TAG\_subrange\_type** entries that describe the constant values associated with the condition. If any child entry has a **DW\_AT\_type** attribute, that attribute describes a type compatible with the comparison type (according to the source language); otherwise the child’s type is the same as the comparison type.

*For conditional variables with alphanumeric types, COBOL permits a source program to provide ranges of alphanumeric constants in the condition. Normally a subrange type entry does not describe ranges of strings; however, this can be represented using bounds attributes that are references to constant entries describing strings. A subrange type entry may refer to constant entries that are siblings of the subrange type entry.*

## 5.9 Enumeration Type Entries

An “enumeration type” is a scalar that can assume one of a fixed number of symbolic values.

An enumeration type is represented by a debugging information entry with the tag **DW\_TAG\_enumeration\_type**.

If a name has been given to the enumeration type in the source program, then the corresponding enumeration type entry has a **DW\_AT\_name** attribute whose value is a null-terminated string containing the enumeration type name.



## Chapter 5. Type Entries

1 The enumeration type entry may have a `DW_AT_type` attribute which refers to  
2 the underlying data type used to implement the enumeration. The entry also  
3 may have a `DW_AT_byte_size` attribute or `DW_AT_bit_size` attribute, whose  
4 value (see Section 2.21 on page 58) is the amount of storage required to hold an  
5 instance of the enumeration. If no `DW_AT_byte_size` or `DW_AT_bit_size`  
6 attribute is present, the size for holding an instance of the enumeration is given  
7 by the size of the underlying data type.

8 If an enumeration type has type safe semantics such that

- 9 1. Enumerators are contained in the scope of the enumeration type, and/or
- 10 2. Enumerators are not implicitly converted to another type

11 then the enumeration type entry may have a `DW_AT_enum_class` attribute,  
12 which is a `flag`. In a language that offers only one kind of enumeration  
13 declaration, this attribute is not required.

14 *In C or C++, the underlying type will be the appropriate integral type determined by the*  
15 *compiler from the properties of the enumeration literal values. A C++ type declaration*  
16 *written using `enum class` declares a strongly typed enumeration and is represented using*  
17 *`DW_TAG_enumeration_type` in combination with `DW_AT_enum_class`.*

18 Each enumeration literal is represented by a debugging information entry with  
19 the tag `DW_TAG_enumerator`. Each such entry is a child of the enumeration  
20 type entry, and the enumerator entries appear in the same order as the  
21 declarations of the enumeration literals in the source program.

22 Each enumerator entry has a `DW_AT_name` attribute, whose value is a  
23 null-terminated string containing the name of the enumeration literal. Each  
24 enumerator entry also has a `DW_AT_const_value` attribute, whose value is the  
25 actual numeric value of the enumerator as represented on the target system.

26 If the enumeration type occurs as the description of a dimension of an array type,  
27 and the stride for that dimension is different than what would otherwise be  
28 determined, then the enumeration type entry has either a `DW_AT_byte_stride` or  
29 `DW_AT_bit_stride` attribute which specifies the separation between successive  
30 elements along the dimension as described in Section 2.19 on page 56. The value  
31 of the `DW_AT_bit_stride` attribute is interpreted as bits and the value of the  
32 `DW_AT_byte_stride` attribute is interpreted as bytes.



## 5.10 Subroutine Type Entries

*It is possible in C to declare pointers to subroutines that return a value of a specific type. In both C and C++, it is possible to declare pointers to subroutines that not only return a value of a specific type, but accept only arguments of specific types. The type of such pointers would be described with a “pointer to” modifier applied to a user-defined type.*

A subroutine type is represented by a debugging information entry with the tag `DW_TAG_subroutine_type`. If a name has been given to the subroutine type in the source program, then the corresponding subroutine type entry has a `DW_AT_name` attribute whose value is a null-terminated string containing the subroutine type name.

If the subroutine type describes a function that returns a value, then the subroutine type entry has a `DW_AT_type` attribute to denote the type returned by the subroutine. If the types of the arguments are necessary to describe the subroutine type, then the corresponding subroutine type entry owns debugging information entries that describe the arguments. These debugging information entries appear in the order that the corresponding argument types appear in the source program.

*In C there is a difference between the types of functions declared using function prototype style declarations and those declared using non-prototype declarations.*

A subroutine entry declared with a function prototype style declaration may have a `DW_AT_prototyped` attribute, which is a flag.

Each debugging information entry owned by a subroutine type entry corresponds to either a formal parameter or the sequence of unspecified parameters of the subprogram type:

1. A formal parameter of a parameter list (that has a specific type) is represented by a debugging information entry with the tag `DW_TAG_formal_parameter`. Each formal parameter entry has a `DW_AT_type` attribute that refers to the type of the formal parameter.
2. The unspecified parameters of a variable parameter list are represented by a debugging information entry with the tag `DW_TAG_unspecified_parameters`.

*C++ const-volatile qualifiers are encoded as part of the type of the “this”-pointer. C++11 reference and rvalue-reference qualifiers are encoded using the `DW_AT_reference` and `DW_AT_rvalue_reference` attributes, respectively. See also Section 5.7.8 on page 124.*

1 A subroutine type entry may have the [DW\\_AT\\_reference](#) or  
 2 [DW\\_AT\\_rvalue\\_reference](#) attribute to indicate that it describes the type of a  
 3 member function with reference or rvalue-reference semantics, respectively.

## 4 5.11 String Type Entries

5 *A “string” is a sequence of characters that have specific semantics and operations that*  
 6 *distinguish them from arrays of characters. Fortran is one of the languages that has a*  
 7 *string type. Note that “string” in this context refers to a target machine concept, not the*  
 8 *class string as used in this document (except for the name attribute).*

9 A string type is represented by a debugging information entry with the tag  
 10 [DW\\_TAG\\_string\\_type](#). If a name has been given to the string type in the source  
 11 program, then the corresponding string type entry has a [DW\\_AT\\_name](#) attribute  
 12 whose value is a null-terminated string containing the string type name.

13 A string type entry may have a [DW\\_AT\\_type](#) attribute describing how each  
 14 character is encoded and is to be interpreted. The value of this attribute is a  
 15 [reference](#) to a [DW\\_TAG\\_base\\_type](#) base type entry. If the attribute is absent, then  
 16 the character is encoded using the system default.

17 *The Fortran 2003 language standard allows string types that are composed of different*  
 18 *types of (same sized) characters. While there is no standard list of character kinds, the*  
 19 *kinds ASCII (see [DW\\_ATE\\_ASCII](#)), ISO\_10646 (see [DW\\_ATE\\_UCS](#)) and DEFAULT are*  
 20 *defined.*

21 The string type entry may have a [DW\\_AT\\_byte\\_size](#) attribute or  
 22 [DW\\_AT\\_bit\\_size](#) attribute, whose value (see Section 2.21 on page 58) is the  
 23 amount of storage needed to hold a value of the string type.

24 The string type entry may also have a [DW\\_AT\\_string\\_length](#) attribute whose  
 25 value is either (a) a [reference](#) (see Section 2.19) to another debugging information  
 26 entry that provides the value of the length of the string, or (b) a location  
 27 description yielding the location where the length of the string is stored in the  
 28 program. If the [DW\\_AT\\_string\\_length](#) attribute is not present, the size of the  
 29 string is assumed to be the amount of storage that is allocated for the string (as  
 30 specified by the [DW\\_AT\\_byte\\_size](#) or [DW\\_AT\\_bit\\_size](#) attribute).

31 The string type entry may also have a [DW\\_AT\\_string\\_length\\_byte\\_size](#) or  
 32 [DW\\_AT\\_string\\_length\\_bit\\_size](#) attribute, whose value (see Section 2.21 on  
 33 page 58) is the size of the data to be retrieved from the location referenced by the  
 34 [DW\\_AT\\_string\\_length](#) attribute. If no byte or bit size attribute is present, the size  
 35 of the data to be retrieved is the same as the size of an address on the target  
 36 machine.

## Chapter 5. Type Entries

1 Prior to DWARF Version 5, the meaning of a *DW\_AT\_byte\_size* attribute depended on  
2 the presence of the *DW\_AT\_string\_length* attribute:

- 3 • If *DW\_AT\_string\_length* was present, *DW\_AT\_byte\_size* specified the size of the  
4 length data to be retrieved from the location specified by the  
5 *DW\_AT\_string\_length* attribute.
- 6 • If *DW\_AT\_string\_length* was not present, *DW\_AT\_byte\_size* specified the  
7 amount of storage allocated for objects of the string type.

8 In DWARF Version 5, *DW\_AT\_byte\_size* always specifies the amount of storage  
9 allocated for objects of the string type.

### 10 5.12 Set Type Entries

11 Pascal provides the concept of a “set,” which represents a group of values of ordinal type.

12 A set is represented by a debugging information entry with the tag  
13 *DW\_TAG\_set\_type*. If a name has been given to the set type, then the set type  
14 entry has a *DW\_AT\_name* attribute whose value is a null-terminated string  
15 containing the set type name.

16 The set type entry has a *DW\_AT\_type* attribute to denote the type of an element  
17 of the set.

18 If the amount of storage allocated to hold each element of an object of the given  
19 set type is different from the amount of storage that is normally allocated to hold  
20 an individual object of the indicated element type, then the set type entry has  
21 either a *DW\_AT\_byte\_size* attribute, or *DW\_AT\_bit\_size* attribute whose value  
22 (see Section 2.21 on page 58) is the amount of storage needed to hold a value of  
23 the set type.

### 24 5.13 Subrange Type Entries

25 Several languages support the concept of a “subrange” type. Objects of the subrange type  
26 can represent only a contiguous subset (range) of values from the type on which the  
27 subrange is defined. Subrange types may also be used to represent the bounds of array  
28 dimensions.

29 A subrange type is represented by a debugging information entry with the tag  
30 *DW\_TAG\_subrange\_type*. If a name has been given to the subrange type, then  
31 the subrange type entry has a *DW\_AT\_name* attribute whose value is a  
32 null-terminated string containing the subrange type name.

## Chapter 5. Type Entries

1 The tag `DW_TAG_generic_subrange` is used to describe arrays with a dynamic  
2 rank. See Section 5.5 on page 115.

3 The subrange entry may have a `DW_AT_type` attribute to describe the type of  
4 object, called the basis type, of whose values this subrange is a subset.

5 If the amount of storage allocated to hold each element of an object of the given  
6 subrange type is different from the amount of storage that is normally allocated  
7 to hold an individual object of the indicated element type, then the subrange  
8 type entry has a `DW_AT_byte_size` attribute or `DW_AT_bit_size` attribute, whose  
9 value (see Section 2.19 on page 56) is the amount of storage needed to hold a  
10 value of the subrange type.

11 The subrange entry may have a `DW_AT_threads_scaled` attribute, which is a  
12 **flag**. If present, this attribute indicates whether this subrange represents a UPC  
13 array bound which is scaled by the runtime `THREADS` value (the number of UPC  
14 threads in this execution of the program).

15 *This allows the representation of a UPC shared array such as*

```
int shared foo[34*THREADS][10][20];
```

16 The subrange entry may have the attributes `DW_AT_lower_bound` and  
17 `DW_AT_upper_bound` to specify, respectively, the lower and upper bound  
18 values of the subrange. The `DW_AT_upper_bound` attribute may be replaced by  
19 a `DW_AT_count` attribute, whose value describes the number of elements in the  
20 subrange rather than the value of the last element. The value of each of these  
21 attributes is determined as described in Section 2.19 on page 56.

22 If the lower bound value is missing, the value is assumed to be a  
23 language-dependent default constant as defined in Table 7.17 on page 242.

24 If the upper bound and count are missing, then the upper bound value is  
25 *unknown*.

26 If the subrange entry has no type attribute describing the basis type, the basis  
27 type is determined as follows:

- 28 1. If there is a lower bound attribute that references an object, the basis type is  
29 assumed to be the same as the type of that object.
- 30 2. Otherwise, if there is an upper bound or count attribute that references an  
31 object, the basis type is assumed to be the same as the type of that object.
- 32 3. Otherwise, the type is assumed to be the same type, in the source language of  
33 the compilation unit containing the subrange entry, as a signed integer with  
34 the same size as an address on the target machine.

1 If the subrange type occurs as the description of a dimension of an array type,  
2 and the stride for that dimension is different than what would otherwise be  
3 determined, then the subrange type entry has either a `DW_AT_byte_stride` or  
4 `DW_AT_bit_stride` attribute which specifies the separation between successive  
5 elements along the dimension as described in Section 2.21 on page 58.

6 *Note that the stride can be negative.*

## 7 **5.14 Pointer to Member Type Entries**

8 *In C++, a pointer to a data or function member of a class or structure is a unique type.*

9 A debugging information entry representing the type of an object that is a pointer  
10 to a structure or class member has the tag `DW_TAG_ptr_to_member_type`.

11 If the pointer to member type has a name, the pointer to member entry has a  
12 `DW_AT_name` attribute, whose value is a null-terminated string containing the  
13 type name.

14 The pointer to member entry has a `DW_AT_type` attribute to describe the type of  
15 the class or structure member to which objects of this type may point.

16 The entry also has a `DW_AT_containing_type` attribute, whose value is a  
17 `reference` to a debugging information entry for the class or structure to whose  
18 members objects of this type may point.

19 The pointer to member entry has a `DW_AT_use_location` attribute whose value  
20 is a location description that computes the address of the member of the class to  
21 which the pointer to member entry points.

22 *The method used to find the address of a given member of a class or structure is common  
23 to any instance of that class or structure and to any instance of the pointer or member  
24 type. The method is thus associated with the type entry, rather than with each instance of  
25 the type.*

26 The `DW_AT_use_location` description is used in conjunction with the location  
27 descriptions for a particular object of the given pointer to member type and for a  
28 particular structure or class instance. The `DW_AT_use_location` attribute expects  
29 two values to be pushed onto the DWARF expression stack before the  
30 `DW_AT_use_location` description is evaluated. The first value pushed is the  
31 value of the pointer to member object itself. The second value pushed is the base  
32 address of the entire structure or union instance containing the member whose  
33 address is being calculated.

1 For an expression such as

```
object.*mbr_ptr
```

2 where *mbr\_ptr* has some pointer to member type, a debugger should:

- 3 1. Push the value of *mbr\_ptr* onto the DWARF expression stack.
- 4 2. Push the base address of *object* onto the DWARF expression stack.
- 5 3. Evaluate the [DW\\_AT\\_use\\_location](#) description given in the type of *mbr\_ptr*.

## 6 5.15 File Type Entries

7 Some languages, such as Pascal, provide a data type to represent files.

8 A file type is represented by a debugging information entry with the tag  
9 [DW\\_TAG\\_file\\_type](#). If the file type has a name, the file type entry has a  
10 [DW\\_AT\\_name](#) attribute, whose value is a null-terminated string containing the  
11 type name.

12 The file type entry has a [DW\\_AT\\_type](#) attribute describing the type of the objects  
13 contained in the file.

14 The file type entry also has a [DW\\_AT\\_byte\\_size](#) or [DW\\_AT\\_bit\\_size](#) attribute,  
15 whose value (see Section 2.19 on page 56) is the amount of storage need to hold a  
16 value of the file type.

## 17 5.16 Dynamic Type Entries

18 Some languages such as Fortran 90, provide types whose values may be dynamically  
19 allocated or associated with a variable under explicit program control. However, unlike  
20 the pointer type in C or C++, the indirection involved in accessing the value of the  
21 variable is generally implicit, that is, not indicated as part of the program source.

22 A dynamic type entry is used to declare a dynamic type that is “just like” another  
23 non-dynamic type without needing to replicate the full description of that other  
24 type.

25 A dynamic type is represented by a debugging information entry with the tag  
26 [DW\\_TAG\\_dynamic\\_type](#). If a name has been given to the dynamic type, then the  
27 dynamic type has a [DW\\_AT\\_name](#) attribute whose value is a null-terminated  
28 string containing the dynamic type name.



1 A dynamic type entry has a `DW_AT_type` attribute whose value is a reference to  
2 the type of the entities that are dynamically allocated.

3 A dynamic type entry also has a `DW_AT_data_location`, and may also have  
4 `DW_AT_allocated` and/or `DW_AT_associated` attributes as described in Section  
5 5.18. A `DW_AT_data_location`, `DW_AT_allocated` or `DW_AT_associated`  
6 attribute may not occur on a dynamic type entry if the same kind of attribute  
7 already occurs on the type referenced by the `DW_AT_type` attribute.

## 8 5.17 Template Alias Entries

9 *In C++, a template alias is a form of typedef that has template parameters. DWARF does*  
10 *not represent the template alias definition but does represent instantiations of the alias.*

11 A type named using a template alias is represented by a debugging information  
12 entry with the tag `DW_TAG_template_alias`. The template alias entry has a  
13 `DW_AT_name` attribute whose value is a null-terminated string containing the  
14 name of the template alias. The template alias entry has child entries describing  
15 the template actual parameters (see Section 2.23 on page 58).

## 16 5.18 Dynamic Properties of Types

17 *The `DW_AT_data_location`, `DW_AT_allocated` and `DW_AT_associated` attributes*  
18 *described in this section are motivated for use with `DW_TAG_dynamic_type` entries but*  
19 *can be used for any other type as well.*

### 20 5.18.1 Data Location

21 *Some languages may represent objects using descriptors to hold information, including a*  
22 *location and/or run-time parameters, about the data that represents the value for that*  
23 *object.*

24 The `DW_AT_data_location` attribute may be used with any type that provides  
25 one or more levels of hidden indirection and/or run-time parameters in its  
26 representation. Its value is a location description. The result of evaluating this  
27 description yields the location of the data for an object. When this attribute is  
28 omitted, the address of the data is the same as the address of the object.

1 *This location description will typically begin with `DW_OP_push_object_address` which*  
2 *loads the address of the object which can then serve as a descriptor in subsequent*  
3 *calculation. For an example using `DW_AT_data_location` for a Fortran 90 array, see*  
4 *Appendix [D.2.1](#) on page 304.*

## 5 5.18.2 Allocation and Association Status

6 *Some languages, such as Fortran 90, provide types whose values may be dynamically*  
7 *allocated or associated with a variable under explicit program control.*

8 The `DW_AT_allocated` attribute may be used with any type for which objects of  
9 the type can be explicitly allocated and deallocated. The presence of the attribute  
10 indicates that objects of the type are allocatable and deallocatable. The integer  
11 value of the attribute (see below) specifies whether an object of the type is  
12 currently allocated or not.

13 The `DW_AT_associated` attribute may optionally be used with any type for  
14 which objects of the type can be dynamically associated with other objects. The  
15 presence of the attribute indicates that objects of the type can be associated. The  
16 integer value of the attribute (see below) indicates whether an object of the type  
17 is currently associated or not.

18 The value of these attributes is determined as described in Section [2.19](#) on  
19 [page 56](#). A non-zero value is interpreted as allocated or associated, and zero is  
20 interpreted as not allocated or not associated.

21 *For Fortran 90, if the `DW_AT_associated` attribute is present, the type has the*  
22 *POINTER property where either the parent variable is never associated with a dynamic*  
23 *object or the implementation does not track whether the associated object is static or*  
24 *dynamic. If the `DW_AT_allocated` attribute is present and the `DW_AT_associated`*  
25 *attribute is not, the type has the ALLOCATABLE property. If both attributes are present,*  
26 *then the type should be assumed to have the POINTER property (and not*  
27 *ALLOCATABLE); the `DW_AT_allocated` attribute may then be used to indicate that the*  
28 *association status of the object resulted from execution of an ALLOCATE statement*  
29 *rather than pointer assignment.*

30 *For examples using `DW_AT_allocated` for Ada and Fortran 90 arrays, see Appendix [D.2](#)*  
31 *on page 304.*



### 5.18.3 Array Rank

*The Fortran language supports “assumed-rank arrays”. The rank (the number of dimensions) of an assumed-rank array is unknown at compile time. The Fortran runtime stores the rank in an array descriptor.*

The presence of the attribute indicates that an array’s rank (number of dimensions) is dynamic, and therefore unknown at compile time. The value of the `DW_AT_rank` attribute is either an integer constant or a DWARF expression whose evaluation yields the dynamic rank.

The bounds of an array with dynamic rank are described using a `DW_TAG_generic_subrange` entry, which is the dynamic rank array equivalent of `DW_TAG_subrange_type`. The difference is that a `DW_TAG_generic_subrange` entry contains generic lower/upper bound and stride expressions that need to be evaluated for each dimension. Before any expression contained in a `DW_TAG_generic_subrange` can be evaluated, the dimension for which the expression is to be evaluated needs to be pushed onto the stack. The expression will use it to find the offset of the respective field in the array descriptor metadata.

*A producer is free to choose any layout for the array descriptor. In particular, the upper and lower bounds and stride values do not need to be bundled into a structure or record, but could be laid end to end in the containing descriptor, pointed to by the descriptor, or even allocated independently of the descriptor.*

Dimensions are enumerated 0 to  $rank - 1$  in source program order.

*For an example in Fortran 2008, see Section D.2.3 on page 313.*

## Chapter 5. Type Entries

*(empty page)*

# Chapter 6

## Other Debugging Information

This section describes debugging information that is not represented in the form of debugging information entries and is not contained within a `.debug_info` section.

In the descriptions that follow, these terms are used to specify the representation of DWARF sections:

- initial length, section offset and section length, which are defined in Sections [7.2.2 on page 192](#) and [7.4 on page 204](#).
- sbyte, ubyte, uhalf and uword, which are defined in Section [7.31 on page 257](#).
- MBZ, which indicates that a value or the contents of a field must be zero.

### 6.1 Accelerated Access

*A debugger frequently needs to find the debugging information for a program entity defined outside of the compilation unit where the debugged program is currently stopped. Sometimes the debugger will know only the name of the entity; sometimes only the address. To find the debugging information associated with a global entity by name, using the DWARF debugging information entries alone, a debugger would need to run through all entries at the highest scope within each compilation unit.*

*Similarly, in languages in which the name of a type is required to always refer to the same concrete type (such as C++), a compiler may choose to elide type definitions in all compilation units except one. In this case a debugger needs a rapid way of locating the concrete type definition by name. As with the definition of global data objects, this would require a search of all the top level type definitions of all compilation units in a program.*

## Chapter 6. Other Debugging Information

1 *To find the debugging information associated with a subroutine, given an address, a*  
2 *debugger can use the low and high PC attributes of the compilation unit entries to*  
3 *quickly narrow down the search, but these attributes only cover the range of addresses for*  
4 *the text associated with a compilation unit entry. To find the debugging information*  
5 *associated with a data object, given an address, an exhaustive search would be needed.*  
6 *Furthermore, any search through debugging information entries for different compilation*  
7 *units within a large program would potentially require the access of many memory pages,*  
8 *probably hurting debugger performance.*

9 To make lookups of program entities (including data objects, functions and  
10 types) by name or by address faster, a producer of DWARF information may  
11 provide two different types of tables containing information about the  
12 debugging information entries owned by a particular compilation unit entry in a  
13 more condensed format.

### 14 **6.1.1 Lookup by Name**

15 For lookup by name, a name index is maintained in a separate object file section  
16 named `.debug_names`.

17 *The `.debug_names` section is new in DWARF Version 5, and supersedes the*  
18 *`.debug_pubnames` and `.debug_pubtypes` sections of earlier DWARF versions. While*  
19 *`.debug_names` and either `.debug_pubnames` and/or `.debug_pubtypes` sections cannot*  
20 *both occur in the same compilation unit, both may be found in the set of units that make*  
21 *up an executable or shared object.*

22 The index consists primarily of two parts: a list of names, and a list of index  
23 entries. A name, such as a subprogram name, type name, or variable name, may  
24 have several defining declarations in the debugging information. In this case, the  
25 entry for that name in the list of names will refer to a sequence of index entries in  
26 the second part of the table, each corresponding to one defining declaration in  
27 the `.debug_info` section.

28 The name index may also contain an optional hash table for faster lookup.

29 A relocatable object file may contain a "per-CU" index, which provides an index  
30 to the names defined in that compilation unit.

31 An executable or shareable object file may contain either a collection of "per-CU"  
32 indexes, simply copied from each relocatable object file, or the linker may  
33 produce a "per-module" index by combining the per-CU indexes into a single  
34 index that covers the entire module.

1 **6.1.1.1 Contents of the Name Index**

2 The name index must contain an entry for each debugging information entry that  
3 defines a named subprogram, label, variable, type, or namespace, subject to the  
4 following rules:

- 5 • All non-defining declarations (that is, debugging information entries with a  
6 [DW\\_AT\\_declaration](#) attribute) are excluded.
- 7 • [DW\\_TAG\\_namespace](#) debugging information entries without a  
8 [DW\\_AT\\_name](#) attribute are included with the name “(anonymous  
9 namespace)”.
- 10 • All other debugging information entries without a [DW\\_AT\\_name](#) attribute  
11 are excluded.
- 12 • [DW\\_TAG\\_subprogram](#), [DW\\_TAG\\_inlined\\_subroutine](#), and  
13 [DW\\_TAG\\_label](#) debugging information entries without an address  
14 attribute ([DW\\_AT\\_low\\_pc](#), [DW\\_AT\\_high\\_pc](#), [DW\\_AT\\_ranges](#), or  
15 [DW\\_AT\\_entry\\_pc](#)) are excluded.
- 16 • [DW\\_TAG\\_variable](#) debugging information entries with a [DW\\_AT\\_location](#)  
17 attribute that includes a [DW\\_OP\\_addr](#) or [DW\\_OP\\_form\\_tls\\_address](#)  
18 operator are included; otherwise, they are excluded.
- 19 • If a subprogram or inlined subroutine is included, and has a  
20 [DW\\_AT\\_linkage\\_name](#) attribute, there will be an additional index entry for  
21 the linkage name.

22 For the purposes of determining whether a debugging information entry has a  
23 particular attribute (such as [DW\\_AT\\_name](#)), if debugging information entry *A*  
24 has a [DW\\_AT\\_specification](#) or [DW\\_AT\\_abstract\\_origin](#) attribute pointing to  
25 another debugging information entry *B*, any attributes of *B* are considered to be  
26 part of *A*.

27 *The intent of the above rules is to provide the consumer with some assurance that looking*  
28 *up an unqualified name in the index will yield all relevant debugging information entries*  
29 *that provide a defining declaration at global scope for that name.*

30 *A producer may choose to implement additional rules for what names are placed in the*  
31 *index, and may communicate those rules to a cooperating consumer via augmentation*  
32 *sequence as described below.*

1 **6.1.1.2 Structure of the Name Index**

2 Logically, the name index can be viewed as a list of names, with a list of index  
3 entries for each name. Each index entry corresponds to a debugging information  
4 entry that matches the criteria given in the previous section. For example, if one  
5 compilation unit has a function named `fred` and another has a struct named  
6 `fred`, a lookup for “fred” will find the list containing those two index entries.

7 The index section contains eight individual parts, as illustrated in Figure 6.1  
8 following.

- 9 1. A header, describing the layout of the section.
- 10 2. A list of compile units (CUs) referenced by this index.
- 11 3. A list of local type units (TUs) referenced by this index that are present in  
12 this object file.
- 13 4. A list of foreign type units (TUs) referenced by this index that are not  
14 present in this object file (that is, that have been placed in a split DWARF  
15 object file as described in 7.3.2 on page 195).
- 16 5. An optional hash lookup table.
- 17 6. The name table.
- 18 7. An abbreviations table, similar to the one used by the `.debug_info` section.
- 19 8. The entry pool, containing a list of index entries for each name in the name  
20 list.

21 The formats of the header and the hash lookup table are described in Section  
22 6.1.1.4 on page 149.

23 The list of CUs and the list of local TUs are each an array of offsets, each of which  
24 is the offset of a compile unit or a type unit in the `.debug_info` section. For a  
25 per-CU index, there is a single CU entry, and there may be a TU entry for each  
26 type unit generated in the same translation unit as the single CU. For a  
27 per-module index, there will be one CU entry for each compile unit in the  
28 module, and one TU entry for each unique type unit in the module. Each list is  
29 indexed starting at 0.

30 The list of foreign TUs is an array of 64-bit (`DW_FORM_ref_sig8`) type  
31 signatures, representing types referenced by the index whose definitions have  
32 been placed in a different object file (that is, a split DWARF object). This list may  
33 be empty. The foreign TU list immediately follows the local TU list and they both  
34 use the same index, so that if there are  $N$  local TU entries, the index for the first  
35 foreign TU is  $N$ .

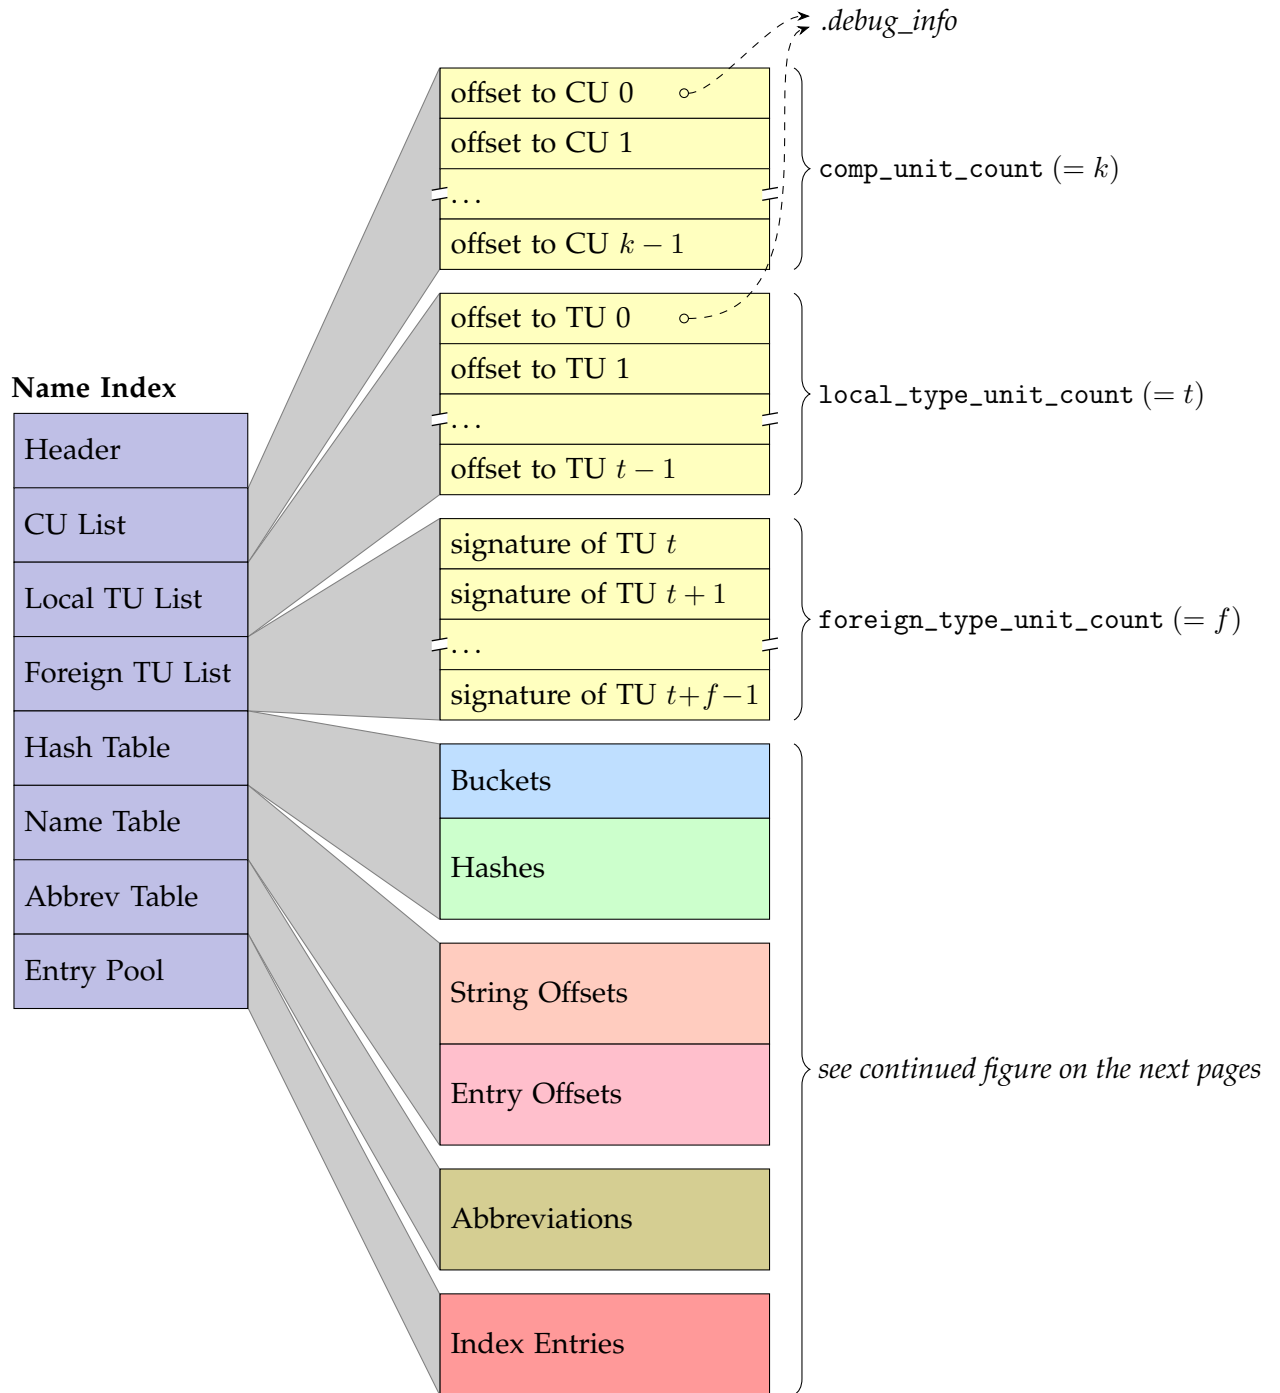
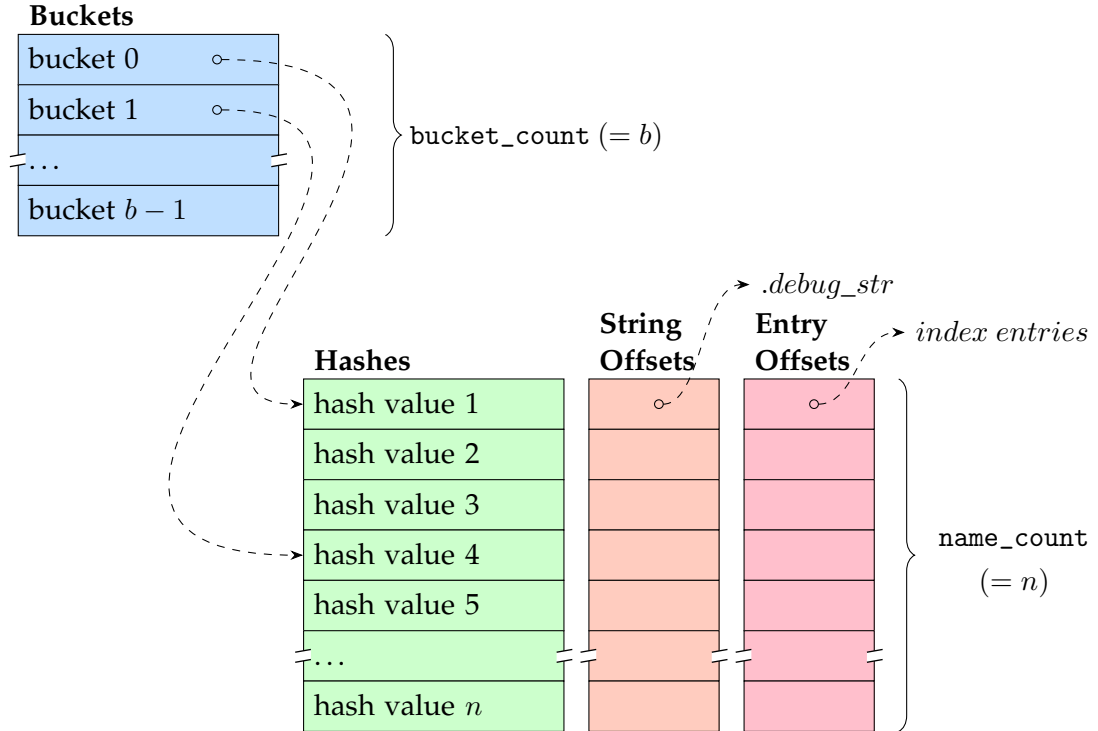


Figure 6.1: Name Index Layout



**Abbreviations**

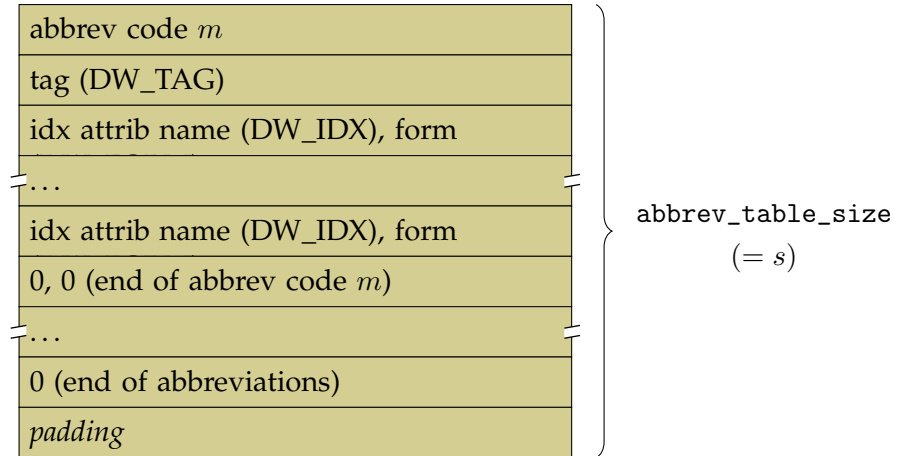


Figure 6.1: Name Index Layout (continued)



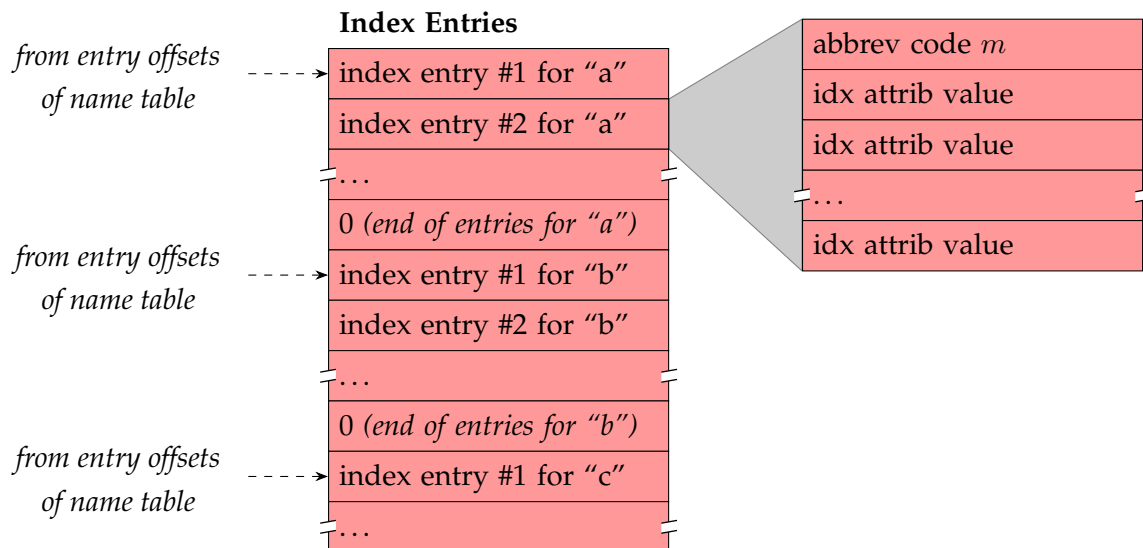


Figure 6.1: Name Index Layout (concluded)

1 The name table is logically a table with a row for each unique name in the index,  
 2 and two columns. The first column contains a reference to the name, as a string.  
 3 The second column contains the offset within the entry pool of the list of index  
 4 entries for the name.

5 The abbreviations table describes the formats of the entries in the entry pool.  
 6 Like the DWARF abbreviations table in the `.debug_abbrev` section, it defines one  
 7 or more abbreviation codes. Each abbreviation code provides a DWARF tag  
 8 value followed by a list of pairs that defines an attribute and form code used by  
 9 entries with that abbreviation code.

10 The entry pool contains all the index entries, grouped by name. The second  
 11 column of the name list points to the first index entry for the name, and all the  
 12 index entries for that name are placed one after the other.

13 Each index entry begins with an unsigned LEB128 abbreviation code. The  
 14 abbreviation list for that code provides the DWARF tag value for the entry as  
 15 well as the set of attributes provided by the entry and their forms.

16 The standard index attributes (see Table 6.1 on page 153) are:

- 17 • Compilation Unit (CU), a reference to an entry in the list of CUs. In a  
 18 per-CU index, index entries without this index attribute implicitly refer to  
 19 the single CU.

## Chapter 6. Other Debugging Information

- 1 • Type Unit (TU), a reference to an entry in the list of local or foreign TUs.
- 2 • Debugging information entry offset within the CU or TU.
- 3 • Parent debugging information entry, a reference to the index entry for the
- 4 parent. This is represented as the offset of the entry relative to the start of
- 5 the entry pool.
- 6 • Type hash, an 8-byte hash of the type declaration.

7 It is possible that an indexed debugging information entry has a parent that is  
8 not indexed (for example, if its parent does not have a name attribute). In such a  
9 case, a parent index attribute may point to a nameless index entry (that is, one  
10 that cannot be reached from any entry in the name table), or it may point to the  
11 nearest ancestor that does have an index entry.

12 A producer may define additional vendor-specific index attributes, and a  
13 consumer will be able to ignore and skip over any index attributes it is not  
14 prepared to handle.

15 When an index entry refers to a foreign type unit, it may have index attributes  
16 for both CU and (foreign) TU. For such entries, the CU index attribute gives the  
17 consumer a reference to the CU that may be used to locate a split DWARF object  
18 file that contains the type unit.

19 *The type hash index attribute, not to be confused with the type signature for a TU, may*  
20 *be provided for type entries whose declarations are not in a type unit, for the convenience*  
21 *of link-time or post-link utilities that wish to de-duplicate type declarations across*  
22 *compilation units. The type hash, however, is computed by the same method as specified*  
23 *for type signatures.*

24 The last entry for each name is followed by a zero byte that terminates the list.  
25 There may be gaps between the lists.

### 26 6.1.1.3 Per-CU versus Per-Module Indexes

27 *In a per-CU index, the CU list may have only a single entry, and index entries may omit*  
28 *the CU attribute. (Cross-module or link-time optimization, however, may produce an*  
29 *object file with several compile units in one object. A compiler in this case may produce a*  
30 *separate index for each CU, or a combined index for all CUs. In the latter case, index*  
31 *entries will require the CU attribute.) Most name table entries may have only a single*  
32 *index entry for each, but sometimes a name may be used in more than one context and*  
33 *will require multiple index entries, each pointing to a different debugging information*  
34 *entry.*

## Chapter 6. Other Debugging Information

1 *When linking object files containing per-CU indexes, the linker may choose to*  
2 *concatenate the indexes as ordinary sections, or it may choose to combine the input*  
3 *indexes into a single per-module index.*

4 *A per-module index will contain a number of CUs, and each index entry contains a CU*  
5 *attribute or a TU attribute to identify which CU or TU contains the debugging*  
6 *information entry being indexed. When a given name is used in multiple CUs or TUs, it*  
7 *will typically have a series of index entries pointing to each CU or TU where it is*  
8 *declared. For example, an index entry for a C++ namespace needs to list each occurrence,*  
9 *since each CU may contribute additional names to the namespace, and the consumer*  
10 *needs to find them all. On the other hand, some index entries do not need to list more*  
11 *than one definition; for example, with the one-definition rule in C++, duplicate entries for*  
12 *a function may be omitted, since the consumer only needs to find one declaration.*  
13 *Likewise, a per-module index needs to list only a single copy of a type declaration*  
14 *contained in a type unit.*

15 *For the benefit of link-time or post-link utilities that consume per-CU indexes and*  
16 *produce a per-module index, the per-CU index entries provide the tag encoding for the*  
17 *original debugging information entry, and may provide a type hash for certain types that*  
18 *may benefit from de-duplication. For example, the standard declaration of the typedef*  
19 *wint32\_t is likely to occur in many CUs, but a combined per-module index needs to*  
20 *retain only one; a user declaration of a typedef mytype may refer to a different type at*  
21 *each occurrence, and a combined per-module index retains each unique declaration of that*  
22 *type.*

### 23 **6.1.1.4 Data Representation of the Name Index**

24 The name index is placed in a section named `.debug_names`, and consists of the  
25 eight parts described in the following sections.

#### 26 **6.1.1.4.1 Section Header**

27 The section header contains the following fields:

- 28 1. `unit_length` (initial length)  
29 The length of this contribution to the name index section, not including the  
30 length field itself (see Section 7.2.2 on page 192).
- 31 2. `version` (uhalf)  
32 A version number (see Section 7.19 on page 245). This number is specific to  
33 the name index table and is independent of the DWARF version number.
- 34 3. `padding` (uhalf)  
35 Reserved to DWARF (must be zero).

## Chapter 6. Other Debugging Information

- 1 4. `comp_unit_count` (uword)  
2 The number of CUs in the CU list.
- 3 5. `local_type_unit_count` (uword)  
4 The number of TUs in the local TU list.
- 5 6. `foreign_type_unit_count` (uword)  
6 The number of TUs in the foreign TU list.
- 7 7. `bucket_count` (uword)  
8 The number of hash buckets in the hash lookup table. If there is no hash  
9 lookup table, this field contains 0.
- 10 8. `name_count` (uword)  
11 The number of unique names in the index.
- 12 9. `abbrev_table_size` (uword)  
13 The size in bytes of the abbreviations table.
- 14 10. `augmentation_size` (uword)  
15 The size in bytes of the augmentation sequence. This value must be a  
16 multiple of four.
- 17 11. `augmentation` (sequence of ubyte)  
18 A vendor-specific sequence of bytes, which provides additional information  
19 about the contents of this index. If provided, the sequence begins with four  
20 bytes which serve as a vendor ID. The remainder of the sequence is meant to  
21 be read by a cooperating consumer, and its contents and interpretation are  
22 not specified here. The block is padded with zero bytes to a multiple of four  
23 bytes in length.  
  
24 *The presence of an unrecognized augmentation vendor ID does not make it*  
25 *impossible for a consumer to process data in the `.debug_names` section. The*  
26 *augmentation sequence only provides hints to the consumer regarding the*  
27 *completeness of the set of names in the index.*

### 28 **6.1.1.4.2 List of CUs**

29 The list of CUs immediately follows the header. Each entry in the list is an offset  
30 of the corresponding compilation unit in the `.debug_info` section. In the  
31 DWARF-32 format, a section offset is 4 bytes, while in the DWARF-64 format, a  
32 section offset is 8 bytes.

33 The total number of entries in the list is given by `comp_unit_count`. There must  
34 be at least one CU.

1     **6.1.1.4.3 List of Local TUs**

2     The list of local TUs immediately follows the list of CUs. Each entry in the list is  
3     an offset of the corresponding type unit in the `.debug_info` section. In the  
4     DWARF-32 format, a section offset is 4 bytes, while in the DWARF-64 format, a  
5     section offset is 8 bytes.

6     Any local TU entry with a maximum representable value is considered not  
7     present. Any index entry referencing such a local TU entry should be ignored.

8     The total number of entries in the list is given by `local_type_unit_count`. This  
9     list may be empty.

10    **6.1.1.4.4 List of Foreign TUs**

11    The list of foreign TUs immediately follows the list of local TUs. Each entry in  
12    the list is a 8-byte type signature (as described by [DW\\_FORM\\_ref\\_sig8](#)).

13    The number of entries in the list is given by `foreign_type_unit_count`. This list  
14    may be empty.

15    **6.1.1.4.5 Hash Lookup Table**

16    The optional hash lookup table immediately follows the list of type signatures.

17    The hash lookup table is actually two separate arrays: an array of buckets,  
18    followed immediately by an array of hashes. The number of entries in the  
19    buckets array is given by `bucket_count`, and the number of entries in the hashes  
20    array is given by `name_count`. Each array contains 4-byte unsigned integers.

21    Symbols are entered into the hash table by first computing a hash value from the  
22    symbol name. The hash is computed using the "DJB" hash function described in  
23    Section [7.33 on page 262](#). Given a hash value for the symbol, the symbol is  
24    entered into a bucket whose index is the hash value modulo `bucket_count`. The  
25    buckets array is indexed starting at 0.

26    For the purposes of the hash computation, each symbol name should be folded  
27    according to the simple case folding algorithm defined in the "Caseless  
28    Matching" subsection of Section 5.18 ("Case Mappings") of the Unicode Standard,  
29    Version 9.0.0. The original symbol name, as it appears in the source code, should  
30    be stored in the name table.

31    *Thus, two symbols that differ only by case will hash to the same slot, but the consumer*  
32    *will be able to distinguish the names when appropriate.*

33    The simple case folding algorithm is further described in the `CaseFolding.txt` file  
34    distributed with the Unicode Character Database. That file defines four classes of  
35    mappings: Common (C), Simple (S), Full (F), and Turkish (T). The hash

## Chapter 6. Other Debugging Information

1 computation specified here uses the C + S mappings only, which do not affect the  
2 total length of the string, with the addition that Turkish upper case dotted 'İ' and  
3 lower case dotless 'ı' are folded to the Latin lower case 'i'.

4 Each bucket contains the index of an entry in the hashes array. The hashes array  
5 is indexed starting at 1, and an empty bucket is represented by the value 0.

6 The hashes array contains a sequence of the full hash values for each symbol. All  
7 symbols that have the same index into the bucket list follow one another in the  
8 hashes array, and the indexed entry in the bucket list refers to the first symbol.  
9 When searching for a symbol, the search starts at the index given by the bucket,  
10 and continues either until a matching symbol is found or until a hash value from  
11 a different bucket is found. If two different symbol names produce the same hash  
12 value, that hash value will occur twice in the hashes array. Thus, if a matching  
13 hash value is found, but the name does not match, the search continues visiting  
14 subsequent entries in the hashes table.

15 When a matching hash value is found in the hashes array, the index of that entry  
16 in the hashes array is used to find the corresponding entry in the name table.

### 17 **6.1.1.4.6 Name Table**

18 The name table immediately follows the hash lookup table. It consists of two  
19 arrays: an array of string offsets, followed immediately by an array of entry  
20 offsets. The items in both arrays are section offsets: 4-byte unsigned integers for  
21 the DWARF-32 format or 8-byte unsigned integers for the DWARF-64 format.  
22 The string offsets in the first array refer to names in the `.debug_str` (or  
23 `.debug_str.dwo`) section. The entry offsets in the second array refer to index  
24 entries, and are relative to the start of the entry pool area.

25 These two arrays are indexed starting at 1, and correspond one-to-one with each  
26 other. The length of each array is given by `name_count`.

27 If there is a hash lookup table, the hashes array corresponds on a one-to-one  
28 basis with the string offsets array and with the entry offsets array.

29 *If there is no hash lookup table, there is no ordering requirement for the name table.*

**6.1.1.4.7 Abbreviations Table**

The abbreviations table immediately follows the name table. This table consists of a series of abbreviation declarations. Its size is given by `abbrev_table_size`.

Each abbreviation declaration defines the tag and other attributes for a particular form of index entry. Each declaration starts with an unsigned LEB128 number representing the abbreviation code itself. It is this code that appears at the beginning of an index entry. The abbreviation code must not be 0.

The abbreviation code is followed by another unsigned LEB128 number that encodes the tag of the debugging information entry corresponding to the index entry.

Following the tag encoding is a series of attribute specifications. Each index attribute consists of two parts: an unsigned LEB128 number that represents the index attribute, and another unsigned LEB128 number that represents the index attribute's form (as described in Section 7.5.4 on page 218). The series of attribute specifications ends with an entry containing 0 for the attribute and 0 for the form.

The index attributes and their meanings are listed in Table 6.1.

Table 6.1: Index attribute encodings

Index attribute name	Meaning
<code>DW_IDX_compile_unit</code>	Index of CU
<code>DW_IDX_type_unit</code>	Index of TU (local or foreign)
<code>DW_IDX_die_offset</code>	Offset of DIE within CU or TU
<code>DW_IDX_parent</code>	Index of name table entry for parent
<code>DW_IDX_type_hash</code>	Hash of type declaration
<code>DW_IDX_external</code>	Whether <code>DW_AT_external</code> is present on the declaration (flag)

The abbreviations table ends with an entry consisting of a single 0 byte for the abbreviation code. The size of the table given by `abbrev_table_size` may include optional padding following the terminating 0 byte.

**6.1.1.4.8 Entry Pool**

The entry pool immediately follows the abbreviations table. Each entry in the entry offsets array in the name table (see Section 6.1.1.4.6) points to an offset in the entry pool, where a series of index entries for that name is located.

1 Each index entry in the series begins with an abbreviation code, and is followed  
2 by the index attribute values described by the abbreviation declaration for that  
3 code. The last index entry in the series is followed by a terminating entry whose  
4 abbreviation code is 0.

5 Each index entry has a flag indicating whether the corresponding DIE has the  
6 `DW_AT_external` attribute with a true value. If the `DW_IDX_external` attribute is  
7 missing from an entry, it means that `DW_AT_external` is false for that DIE.

8 Gaps are not allowed between entries in a series (that is, the entries for a single  
9 name must all be contiguous), but there may be gaps between series.

10 *For example, a producer/consumer combination may find it useful to maintain alignment.*

11 The size of the entry pool is the remaining size of the contribution to the index  
12 section, as defined by the `unit_length` header field.

### 13 6.1.2 Lookup by Address

14 For lookup by address, a table is maintained in a separate object file section  
15 called `.debug_aranges`. The table consists of sets of variable length entries, each  
16 set describing the portion of the program's address space that is covered by a  
17 single compilation unit.

18 Each set begins with a header containing five values:

- 19 1. `unit_length` (initial length)  
20 The length of this contribution to the address lookup section, not including  
21 the length field itself (see Section 7.2.2 on page 192).
- 22 2. `version` (uhalf)  
23 A version number (see Section 7.21 on page 246). This number is specific to  
24 the address lookup table and is independent of the DWARF version number.
- 25 3. `debug_info_offset` (section offset)  
26 The offset from the beginning of the `.debug_info` section of the compilation  
27 unit header referenced by the set.
- 28 4. `address_size` (ubyte)  
29 The size of an address in bytes on the target architecture.
- 30 5. `reserved`<sup>1</sup> (ubyte, MBZ)

31 This header is followed by a variable number of address range descriptors. Each  
32 descriptor is a pair consisting of the beginning address of a range of text or data

---

<sup>1</sup>This allows backward compatible support of the deprecated `segment_selector_size` field which was defined in DWARF Version 5 and earlier.



1 covered by some entry owned by the corresponding compilation unit, followed  
2 by the length of that range. A particular set is terminated by an entry consisting  
3 of two zeroes. By scanning the table, a debugger can quickly decide which  
4 compilation unit to look in to find the debugging information for an object that  
5 has a given address.

6 A range description entry whose address is the reserved address (see Section  
7 [2.4.1 on page 26](#)), indicates a non-existent range, which is equivalent to omitting  
8 the range description.

9 *If the range of addresses covered by the text and/or data of a compilation unit is not*  
10 *contiguous, then there may be multiple address range descriptors for that compilation*  
11 *unit.*

## 12 **6.2 Line Number Information**

13 *A source-level debugger needs to know how to associate locations in the source files with*  
14 *the corresponding machine instruction addresses in the executable or the shared object*  
15 *files used by that executable object file. Such an association makes it possible for the*  
16 *debugger user to specify machine instruction addresses in terms of source locations. This*  
17 *is done by specifying the line number and the source file containing the statement. The*  
18 *debugger can also use this information to display locations in terms of the source files and*  
19 *to single step from line to line, or statement to statement.*

20 Line number information generated for a compilation unit is represented in the  
21 `.debug_line` section of an object file, and optionally also in the `.debug_line_str`  
22 section, and is referenced by a corresponding compilation unit debugging  
23 information entry (see Section [3.1.1 on page 62](#)) in the `.debug_info` section.

24 *Some computer architectures employ more than one instruction set (for example, the*  
25 *ARM and MIPS architectures support a 32-bit as well as a 16-bit instruction set).*  
26 *Because the instruction set is a function of the program counter, it is convenient to*  
27 *encode the applicable instruction set in the `.debug_line` section as well.*

28 *If space were not a consideration, the information provided in the `.debug_line` section*  
29 *could be represented as a large matrix, with one row for each instruction in the emitted*  
30 *object code. The matrix would have columns for:*

- 31 • *the source file name*
- 32 • *the source line number*
- 33 • *the source column number*
- 34 • *whether this instruction is the beginning of a source statement*

## Chapter 6. Other Debugging Information

- 1 • *whether this instruction is the beginning of a basic block*
- 2 • *and so on*

3 *Such a matrix, however, would be impractically large. We shrink it with two techniques.*  
4 *First, we delete from the matrix each row whose file, line, source column and*  
5 *discriminator is identical with that of its predecessors. Any deleted row would never be*  
6 *the beginning of a source statement. Second, we design a byte-coded language for a state*  
7 *machine and store a stream of bytes in the object file instead of the matrix. This language*  
8 *can be much more compact than the matrix. To the line number information a consumer*  
9 *must “run” the state machine to generate the matrix for each compilation unit of interest.*  
10 *The concept of an encoded matrix also leaves room for expansion. In the future, columns*  
11 *can be added to the matrix to encode other things that are related to individual*  
12 *instruction addresses.*

### 13 **6.2.1 Definitions**

14 The following terms are used in the description of the line number information  
15 format:

state machine	The hypothetical machine used by a consumer of the line number information to expand the byte-coded instruction stream into a matrix of line number information.
line number program	A series of byte-coded line number information instructions representing one compilation unit.
basic block	A sequence of instructions where only the first instruction may be a branch target and only the last instruction may transfer control. A subprogram invocation is defined to be an exit from a basic block. <i>A basic block does not necessarily correspond to a specific source code construct.</i>
sequence	A series of contiguous target machine instructions. One compilation unit may emit multiple sequences (that is, not all instructions within a compilation unit are assumed to be contiguous).

## 6.2.2 State Machine Registers

The line number information state machine has a number of registers as shown in Table 6.3 following.

Table 6.3: State machine registers

Register name	Meaning
address	The program-counter value corresponding to a machine instruction generated by the compiler.
op_index	An unsigned integer representing the index of an operation within a VLIW instruction. The index of the first operation is 0. For non-VLIW architectures, this register will always be 0.
file	An unsigned integer indicating the identity of the source file corresponding to a machine instruction. Files are numbered beginning at 0.
line	An unsigned integer indicating a source line number. Lines are numbered beginning at 1. The compiler may emit the value 0 in cases where an instruction cannot be attributed to any source line.
column	An unsigned integer indicating a column number within a source line. Columns are numbered beginning at 1. The value 0 is reserved to indicate that a statement begins at the “left edge” of the line.
is_stmt	A boolean indicating that the current instruction is a recommended breakpoint location. A recommended breakpoint location is intended to “represent” a line, a statement and/or a semantically distinct subpart of a statement.
basic_block	A boolean indicating that the current instruction is the beginning of a basic block.
end_sequence	A boolean indicating that the current address is that of the first byte after the end of a sequence of target machine instructions. <code>end_sequence</code> terminates a sequence of lines; therefore other information in the same row is not meaningful.

*Continued on next page*

Register name	Meaning
prologue_end	A boolean indicating that the current address is one (of possibly many) where execution should be suspended for a breakpoint at the entry of a function.
epilogue_begin	A boolean indicating that the current address is one (of possibly many) where execution should be suspended for a breakpoint just prior to the exit of a function.
isa	An unsigned integer whose value encodes the applicable instruction set architecture for the current instruction. <i>The encoding of instruction sets should be shared by all users of a given architecture. It is recommended that this encoding be defined by the ABI authoring committee for each architecture.</i>
discriminator	An unsigned integer identifying the block to which the current instruction belongs. Discriminator values are assigned arbitrarily by the DWARF producer and serve to distinguish among multiple blocks that may all be associated with the same source file, line, and column. Where only one block exists for a given source position, the discriminator value is zero.

1 The address and `op_index` registers, taken together, form an operation pointer  
2 that can reference any individual operation within the instruction stream.

3 At the beginning of each sequence within a line number program, the state of the  
4 registers is as show in Table 6.4 on the following page.

5 *The `isa` value 0 specifies that the instruction set is the architecturally determined default*  
6 *instruction set. This may be fixed by the ABI, or it may be specified by other means, for*  
7 *example, by the object file description.*

### 8 6.2.3 Line Number Program Instructions

9 The state machine instructions in a line number program belong to one of three  
10 categories:

11 1. special opcodes

12 These have a ubyte opcode field and no operands.

## Chapter 6. Other Debugging Information

Table 6.4: Line number program initial state

address	0
op_index	0
file	0
line	1
column	0
is_stmt	determined by default_is_stmt in the line number program header
basic_block	“false”
end_sequence	“false”
prologue_end	“false”
epilogue_begin	“false”
isa	0
discriminator	0

- 1 *Most of the instructions in a line number program are special opcodes.*
- 2 2. standard opcodes
- 3 These have a ubyte opcode field which may be followed by zero or more
- 4 LEB128 operands (except for [DW\\_LNS\\_fixed\\_advance\\_pc](#), see Section 6.2.5.2
- 5 [on page 169](#)). The opcode implies the number of operands and their
- 6 meanings, but the line number program header also specifies the number of
- 7 operands for each standard opcode.
- 8 One standard opcode ([DW\\_LNS\\_extended\\_op](#)) serves as an escape that
- 9 allows additional opcodes without reducing the number of special opcodes.
- 10 3. extended opcodes
- 11 These have a multiple byte format. The first byte is [DW\\_LNS\\_extended\\_op](#).
- 12 The next bytes are an unsigned LEB128 integer giving the number of bytes in
- 13 the instruction itself (this does not include the first [DW\\_LNS\\_extended\\_op](#)
- 14 byte or the size). The remaining bytes are the instruction itself (which begins
- 15 with a ubyte extended opcode).

## 6.2.4 The Line Number Program Header

The optimal encoding of line number information depends to a certain degree upon the architecture of the target machine. The line number program header provides information used by consumers in decoding the line number program instructions for a particular compilation unit and also provides information used throughout the rest of the line number program.

The line number program for each compilation unit begins with a header containing the following fields in order:

1. `unit_length` (initial length)  
The size in bytes of the line number information for this compilation unit, not including the length field itself (see Section 7.2.2 on page 192).
2. `version` (uhalf)  
A version number (see Section 7.22 on page 247). This number is specific to the line number information and is independent of the DWARF version number.
3. `address_size` (ubyte)  
The size of an address in bytes on the target architecture.  
*The `address_size` field supports the common practice of stripping all but the line number sections (`.debug_line` and `.debug_line_str`) from an executable.*
4. `reserved`<sup>2</sup> (ubyte, MBZ)
5. `header_length`  
The number of bytes following the `header_length` field to the beginning of the first byte of the line number program itself. In the 32-bit DWARF format, this is a 4-byte unsigned length; in the 64-bit DWARF format, this field is an 8-byte unsigned length (see Section 7.4 on page 204).
6. `minimum_instruction_length` (ubyte)  
The size in bytes of the smallest target machine instruction. Line number program opcodes that alter the address and `op_index` registers use this and `maximum_operations_per_instruction` in their calculations.

---

<sup>2</sup>This allows backward compatible support of the deprecated `segment_selector_size` field which was defined in DWARF Version 5 and earlier.

## Chapter 6. Other Debugging Information

1 7. `maximum_operations_per_instruction` (ubyte)

2 The maximum number of individual operations that may be encoded in an  
3 instruction. Line number program opcodes that alter the address and  
4 `op_index` registers use this and `minimum_instruction_length` in their  
5 calculations.

6 For non-VLIW architectures, this field is 1, the `op_index` register is always 0,  
7 and the operation pointer is simply the address register.

8 8. `default_is_stmt` (ubyte)

9 The initial value of the `is_stmt` register.

10 *A simple approach to building line number information when machine instructions*  
11 *are emitted in an order corresponding to the source program is to set*  
12 *`default_is_stmt` to “true” and to not change the value of the `is_stmt` register*  
13 *within the line number program. One matrix entry is produced for each line that has*  
14 *code generated for it. The effect is that every entry in the matrix recommends the*  
15 *beginning of each represented line as a breakpoint location. This is the traditional*  
16 *practice for unoptimized code.*

17 *A more sophisticated approach might involve multiple entries in the matrix for a line*  
18 *number; in this case, at least one entry (often but not necessarily only one) specifies a*  
19 *recommended breakpoint location for the line number. `DW_LNS_negate_stmt`*  
20 *opcodes in the line number program control which matrix entries constitute such a*  
21 *recommendation and `default_is_stmt` might be either “true” or “false.” This*  
22 *approach might be used as part of support for debugging optimized code.*

23 9. `line_base` (sbyte)

24 This parameter affects the meaning of the special opcodes. See below.

25 10. `line_range` (ubyte)

26 This parameter affects the meaning of the special opcodes. See below.

27 11. `opcode_base` (ubyte)

28 The number assigned to the first special opcode.

29 *Opcode base is typically one greater than the highest-numbered standard opcode*  
30 *defined for the specified version of the line number information (12 in DWARF*  
31 *Versions 3 through 6, and 9 in Version 2). If `opcode_base` is less than the typical*  
32 *value, then standard opcode numbers greater than or equal to the opcode base are not*  
33 *used in the line number table of this unit (and the codes are treated as special*  
34 *opcodes). If `opcode_base` is greater than the typical value, then the numbers*  
35 *between that of the highest standard opcode and the first special opcode (not*  
36 *inclusive) are used for vendor specific extensions.*

## Chapter 6. Other Debugging Information

1 12. `standard_opcode_lengths` (array of `ubyte`)

2 This array specifies the number of LEB128 operands for each of the standard  
3 opcodes. The first element of the array corresponds to the opcode whose  
4 value is 1, and the last element corresponds to the opcode whose value is  
5 `opcode_base - 1`.

6 *By increasing `opcode_base`, and adding elements to this array, new standard*  
7 *opcodes can be added, while allowing consumers who do not know about these new*  
8 *opcodes to be able to skip them.*

9 *Codes for vendor specific extensions, if any, are described just like standard opcodes.*

10 13. `directory_format_count` (ULEB128)

11 A count of the number of entries in the following `directory_format_table`  
12 field.

13 14. `directory_format_table` (sequence of record format descriptors)

14 A sequence of record format descriptors. Each descriptor consists the  
15 following:

- 16 • A sequence of field descriptors. Each field descriptor consists of a pair of  
17 unsigned LEB128 values: (a) a content type code (see Sections [6.2.4.1 on](#)  
18 [page 164](#) and [6.2.4.2 on page 166](#)), and (b) a form code (using the  
19 attribute form codes).
- 20 • A pair of zero bytes to terminate the descriptor.

21 The line number program numbers the record format descriptors  
22 sequentially, beginning with 0.

23 The format declarations describe the layout of the entries in the `directories`  
24 field, below.

25 15. `directories_count` (ULEB128)

26 A count of the number of entries in the following `directories` field.

27 16. `directories` (sequence of directory entries)

28 A sequence of directory entries. Each entry consists of:

- 29 • A format code (ULEB128), which selects a record format descriptor from  
30 the `directory_format_table` field, above, by its index.
- 31 • A sequence of fields as described by the selected record format  
32 descriptor.



## Chapter 6. Other Debugging Information

1 Each directory entry describes a path that was searched for included source  
2 files in this compilation, including the compilation directory of the  
3 compilation. (The paths include those directories specified by the user for the  
4 compiler to search and those the compiler searches without explicit direction.)

5 The first path entry is the current directory of the compilation; if that entry is  
6 specified using a relative path, it is relative to the location of the linked image  
7 containing the line table entries (assuming the image has not been moved).  
8 Each additional path entry is either a full path name or is relative to the  
9 current directory of the compilation.

10 The line number program assigns a number (index) to each of the directory  
11 entries in order, beginning with 0.

12 *Prior to DWARF Version 5, the current compilation file did not have a specific entry*  
13 *in the `file_names` field. Starting in DWARF Version 5, the current file name has*  
14 *index 0.*

15 *Note that if a `.debug_line_str` section is present, both the compilation unit*  
16 *debugging information entry and the line number header can share a single copy of*  
17 *the current directory name string.*

18 17. `file_name_format_count` (ULEB128)

19 A count of the number of format descriptors in the following  
20 `file_name_format_table` field.

21 18. `file_name_format_table` (sequence of record format descriptors)

22 A sequence of record format descriptors. Each descriptor consists of:

- 23 • A sequence of field descriptors. Each field descriptor consists of a pair of  
24 unsigned LEB128 values: (a) a content type code (see Sections [6.2.4.1 on](#)  
25 [the next page](#) and [6.2.4.2 on page 166](#)), and (b) a form code (using the  
26 attribute form codes).
- 27 • A pair of zero bytes to terminate the descriptor.

28 The line number program numbers the record format descriptors  
29 sequentially, beginning with 0.

30 The format declarations describe the layout of the entries in the `file_names`  
31 field, below.

32 19. `file_names_count` (ULEB128)

33 A count of the number of file name entries in the following `file_names` field.

## Chapter 6. Other Debugging Information

### 20. `file_names` (sequence of file name entries)

A sequence of file name entries. Each entry consists of:

- A format code (ULEB128), which selects a record format descriptor from the `file_name_format_table`, by its index.
- A sequence of fields as described by the selected record format descriptor.

Each file name entry describes a source file that contributes to the line number information for this compilation or is used in other contexts, such as in a declaration coordinate or a macro file inclusion.

The first file name entry is the primary source file, whose file name exactly matches that given in the `DW_AT_name` attribute in the compilation unit debugging information entry.

The line number program references file names in this sequence beginning with 0, and uses those numbers instead of file names in the line number program that follows.

*Prior to DWARF Version 5, the current compilation file name was not represented in the `file_names` field. In DWARF Version 5 and after, the current compilation file name is explicitly present and has index 0. This is needed to support the common practice of stripping all but the line number sections (`.debug_line` and `.debug_line_str`) from an executable.*

*Note that if a `.debug_line_str` section is present, both the compilation unit debugging information entry and the line number header can share a single copy of the current file name string.*

#### 6.2.4.1 Standard Content Descriptions

DWARF-defined content type codes are used to indicate the type of information that is represented in one component of an include directory or file name description. The following type codes are defined.

##### 1. `DW_LNCT_path`

The component is a null-terminated path name string. If the associated form code is `DW_FORM_string`, then the string occurs immediately in the containing `directories` or `file_names` field. If the form code is `DW_FORM_line_strp`, then the string is included in the `.debug_line_str` section; if the form code is `DW_FORM_strp` or `DW_FORM_strp8`, then the string is included in the `.debug_str` section; if the form code is `DW_FORM_strp_sup` or `DW_FORM_strp_sup8`, then the string is included in

## Chapter 6. Other Debugging Information

1 the supplementary string section. In all cases other than `DW_FORM_string`,  
2 the string's offset occurs immediately in the containing directories or  
3 `file_names` field.

4 In the 32-bit DWARF format, the representation of a `DW_FORM_line_strp`  
5 value is a 4-byte unsigned offset; in the 64-bit DWARF format, it is an 8-byte  
6 unsigned offset (see Section 7.4 on page 204).

7 *Note that this use of `DW_FORM_line_strp` is similar to `DW_FORM_strp` but refers*  
8 *to the `.debug_line_str` section, not `.debug_str`. It is needed to support the*  
9 *common practice of stripping all but the line number sections (`.debug_line` and*  
10 *`.debug_line_str`) from an executable.*

11 In a `.debug_line.dwo` section, the forms `DW_FORM_strx`, `DW_FORM_strx1`,  
12 `DW_FORM_strx2`, `DW_FORM_strx3` and `DW_FORM_strx4` may also be  
13 used. These refer into the `.debug_str_offsets.dwo` section (and indirectly  
14 also the `.debug_str.dwo` section) because no `.debug_line_str_offsets.dwo`  
15 or `.debug_line_str.dwo` sections exist or are defined for use in split objects.  
16 (The form `DW_FORM_string` may also be used, but this precludes the  
17 benefits of string sharing.)

### 18 2. `DW_LNCT_directory_index`

19 The unsigned directory index represents an entry in the directories field of  
20 the header. The index is 0 if the file was found in the current directory of the  
21 compilation (hence, the first directory in the directories field), 1 if it was  
22 found in the second directory in the directories field, and so on.

23 This content code is always paired with one of the forms `DW_FORM_data1`,  
24 `DW_FORM_data2` or `DW_FORM_adata`.

25 *The optimal form for a producer to use (which results in the minimum size for the set*  
26 *of `include_index` fields) depends not only on the number of directories in the*  
27 *directories field, but potentially on the order in which those directories are listed and*  
28 *the number of times each is used in the `file_names` field.*

### 29 3. `DW_LNCT_timestamp`

30 `DW_LNCT_timestamp` indicates that the value is the  
31 implementation-defined time of last modification of the file, or 0 if not  
32 available. It is always paired with one of the forms `DW_FORM_adata`,  
33 `DW_FORM_data4`, `DW_FORM_data8` or `DW_FORM_block`.

## Chapter 6. Other Debugging Information

### 4. `DW_LNCT_size`

`DW_LNCT_size` indicates that the value is the unsigned size of the file in bytes, or 0 if not available. It is paired with one of the forms `DW_FORM_udata`, `DW_FORM_data1`, `DW_FORM_data2`, `DW_FORM_data4` or `DW_FORM_data8`.

### 5. `DW_LNCT_MD5`

`DW_LNCT_MD5` indicates that the value is a 16-byte `MD5` digest of the file contents. It is paired with form `DW_FORM_data16`.

### 6. `DW_LNCT_source`

`DW_LNCT_source` specifies a null-terminated UTF-8 string that constitutes the source text for the program. It is paired with the same forms as `DW_LNCT_path`.

*When the source field is present, consumers use the embedded source instead of accessing the source using the file path provided by the `DW_LNCT_path` field.*

*This is useful for programming languages that support runtime compilation and runtime generation of source text. In these cases, the source text does not reside in any permanent file. For example, the OpenCL C language supports runtime compilation.*

### 7. `DW_LNCT_URL`

`DW_LNCT_URL` specifies a null-terminated UTF-8 string that identifies where the source text for the program is found on the Internet. It is paired with the same forms as `DW_LNCT_path`.

*An example that uses this line number header format is found in Appendix D.5.1 on page 339.*

#### 6.2.4.2 Vendor-defined Content Descriptions

Vendor-defined content descriptions may be defined using content type codes in the range `DW_LNCT_lo_user` to `DW_LNCT_hi_user`. Each such code may be combined with one or more forms from the set: `DW_FORM_block`, `DW_FORM_block1`, `DW_FORM_block2`, `DW_FORM_block4`, `DW_FORM_data1`, `DW_FORM_data2`, `DW_FORM_data4`, `DW_FORM_data8`, `DW_FORM_data16`, `DW_FORM_flag`, `DW_FORM_line_strp`, `DW_FORM_sdata`, `DW_FORM_sec_offset`, `DW_FORM_string`, `DW_FORM_strp`, `DW_FORM_strp8`, `DW_FORM_strp_sup`, `DW_FORM_strp_sup8`, `DW_FORM_strx`, `DW_FORM_strx1`, `DW_FORM_strx2`, `DW_FORM_strx3`, `DW_FORM_strx4` and `DW_FORM_udata`.

1 *If a consumer encounters a vendor-defined content type that it does not understand, it*  
2 *should skip the content data as though it were not present.*

## 3 **6.2.5 The Line Number Program**

4 As stated before, the goal of a line number program is to build a matrix  
5 representing one compilation unit, which may have produced multiple  
6 sequences of target machine instructions. Within a sequence, addresses and  
7 operation pointers may only increase. (Line numbers may decrease in cases of  
8 pipeline scheduling or other optimization.)

### 9 **6.2.5.1 Special Opcodes**

10 Each ubyte special opcode has the following effect on the state machine:

- 11 1. Add a signed integer to the `line` register.
- 12 2. Modify the operation pointer by incrementing the address and `op_index`  
13 registers as described below.
- 14 3. Append a row to the matrix using the current values of the state machine  
15 registers.
- 16 4. Set the `basic_block` register to “false.”
- 17 5. Set the `prologue_end` register to “false.”
- 18 6. Set the `epilogue_begin` register to “false.”
- 19 7. Set the `discriminator` register to 0.

20 All of the special opcodes do those same seven things; they differ from one  
21 another only in what values they add to the `line`, `address` and `op_index`  
22 registers.

23 *Instead of assigning a fixed meaning to each special opcode, the line number program*  
24 *uses several parameters in the header to configure the instruction set. There are two*  
25 *reasons for this. First, although the opcode space available for special opcodes ranges from*  
26 *13 through 255, the lower bound may increase if one adds new standard opcodes. Thus,*  
27 *the `opcode_base` field of the line number program header gives the value of the first*  
28 *special opcode. Second, the best choice of special-opcode meanings depends on the target*  
29 *architecture. For example, for a RISC machine where the compiler-generated code*  
30 *interleaves instructions from different lines to schedule the pipeline, it is important to be*  
31 *able to add a negative value to the `line` register to express the fact that a later instruction*  
32 *may have been emitted for an earlier source line. For a machine where pipeline scheduling*  
33 *never occurs, it is advantageous to trade away the ability to decrease the `line` register (a*

## Chapter 6. Other Debugging Information

1 *standard opcode provides an alternate way to decrease the line number) in return for the*  
2 *ability to add larger positive values to the address register. To permit this variety of*  
3 *strategies, the line number program header defines a `line_base` field that specifies the*  
4 *minimum value which a special opcode can add to the line register and a `line_range`*  
5 *field that defines the range of values it can add to the line register.*

6 A special opcode value is chosen based on the amount that needs to be added to  
7 the `line`, `address` and `op_index` registers. The maximum line increment for a  
8 special opcode is the value of the `line_base` field in the header, plus the value of  
9 the `line_range` field, minus 1 (`line base + line range - 1`). If the desired line  
10 increment is greater than the maximum line increment, a standard opcode must  
11 be used instead of a special opcode. The operation advance represents the  
12 number of operations to skip when advancing the operation pointer.

13 The special opcode is then calculated using the following formula:

14 `opcode =`  
15 `(desired line increment - line_base) +`  
16 `(line_range * operation advance) + opcode_base`

17 If the resulting opcode is greater than 255, a standard opcode must be used  
18 instead.

19 *When `maximum_operations_per_instruction` is 1, the operation advance is simply*  
20 *the address increment divided by the `minimum_instruction_length`.*

21 To decode a special opcode, subtract the `opcode_base` from the opcode itself to  
22 give the *adjusted opcode*. The *operation advance* is the result of the adjusted opcode  
23 divided by the `line_range`. The new address and `op_index` values are given by

24 `adjusted opcode = opcode - opcode_base`  
25 `operation advance = adjusted opcode / line_range`  
26  
27 `new address = address +`  
28 `minimum_instruction_length *`  
29 `((op_index + operation advance) / maximum_operations_per_instruction)`  
30  
31 `new op_index =`  
32 `(op_index + operation advance) % maximum_operations_per_instruction`

33 *When the `maximum_operations_per_instruction` field is 1, `op_index` is always 0*  
34 *and these calculations simplify to those given for addresses in DWARF Version 3 and*  
35 *earlier.*

## Chapter 6. Other Debugging Information

1 The amount to increment the line register is the `line_base` plus the result of the  
2 *adjusted opcode* modulo the `line_range`. That is,

3 
$$\text{line increment} = \text{line\_base} + (\text{adjusted opcode} \% \text{line\_range})$$

4 See [Appendix D.5.2 on page 340](#) for an example.

### 5 6.2.5.2 Standard Opcodes

6 The standard opcodes, their applicable operands and the actions performed by  
7 these opcodes are as follows:

8 1. **DW\_LNS\_copy**

9 The `DW_LNS_copy` opcode takes no operands. It appends a row to the  
10 matrix using the current values of the state machine registers. Then it sets the  
11 discriminator register to 0, and sets the `basic_block`, `prologue_end` and  
12 `epilogue_begin` registers to “false.”

13 2. **DW\_LNS\_advance\_pc**

14 The `DW_LNS_advance_pc` opcode takes a single unsigned LEB128 operand  
15 as the operation `advance` and modifies the `address` and `op_index` registers as  
16 specified in [Section 6.2.5.1 on page 167](#).

17 3. **DW\_LNS\_advance\_line**

18 The `DW_LNS_advance_line` opcode takes a single signed LEB128 operand  
19 and adds that value to the `line` register of the state machine.

20 4. **DW\_LNS\_set\_file**

21 The `DW_LNS_set_file` opcode takes a single unsigned LEB128 operand and  
22 stores it in the `file` register of the state machine.

23 5. **DW\_LNS\_set\_column**

24 The `DW_LNS_set_column` opcode takes a single unsigned LEB128 operand  
25 and stores it in the `column` register of the state machine.

26 6. **DW\_LNS\_negate\_stmt**

27 The `DW_LNS_negate_stmt` opcode takes no operands. It sets the `is_stmt`  
28 register of the state machine to the logical negation of its current value.

29 7. **DW\_LNS\_set\_basic\_block**

30 The `DW_LNS_set_basic_block` opcode takes no operands. It sets the  
31 `basic_block` register of the state machine to “true.”

## Chapter 6. Other Debugging Information

### 8. **DW\_LNS\_const\_add\_pc**

The DW\_LNS\_const\_add\_pc opcode takes no operands. It advances the address and op\_index registers by the increments corresponding to special opcode 255.

*When the line number program needs to advance the address by a small amount, it can use a single special opcode, which occupies a single byte. When it needs to advance the address by up to twice the range of the last special opcode, it can use DW\_LNS\_const\_add\_pc followed by a special opcode, for a total of two bytes. Only if it needs to advance the address by more than twice that range will it need to use both DW\_LNS\_advance\_pc and a special opcode, requiring three or more bytes.*

### 9. **DW\_LNS\_fixed\_advance\_pc**

The DW\_LNS\_fixed\_advance\_pc opcode takes a single uhalf (unencoded) operand and adds it to the address register of the state machine and sets the op\_index register to 0. This is the only standard opcode whose operand is **not** a variable length number. It also does **not** multiply the operand by the minimum\_instruction\_length field of the header.

*Some assemblers may not be able emit DW\_LNS\_advance\_pc or special opcodes because they cannot encode LEB128 numbers or judge when the computation of a special opcode overflows and requires the use of DW\_LNS\_advance\_pc. Such assemblers, however, can use DW\_LNS\_fixed\_advance\_pc instead, sacrificing compression.*

### 10. **DW\_LNS\_set\_prologue\_end**

The DW\_LNS\_set\_prologue\_end opcode takes no operands. It sets the prologue\_end register to “true.”

*When a breakpoint is set on entry to a function, it is generally desirable for execution to be suspended, not on the very first instruction of the function, but rather at a point after the function’s frame has been set up, after any language defined local declaration processing has been completed, and before execution of the first statement of the function begins. Debuggers generally cannot properly determine where this point is. This command allows a compiler to communicate the location(s) to use.*

*In the case of optimized code, there may be more than one such location; for example, the code might test for a special case and make a fast exit prior to setting up the frame.*

*Note that the function to which the prologue end applies cannot be directly determined from the line number information alone; the function must be determined in combination with the subroutine information entries of the compilation (including inlined subroutines).*



11. **DW\_LNS\_set\_epilogue\_begin**

The DW\_LNS\_set\_epilogue\_begin opcode takes no operands. It sets the epilogue\_begin register to “true.”

*When a breakpoint is set on the exit of a function or execution steps over the last executable statement of a function, it is generally desirable to suspend execution after completion of the last statement but prior to tearing down the frame (so that local variables can still be examined). Debuggers generally cannot properly determine where this point is. This command allows a compiler to communicate the location(s) to use.*

*Note that the function to which the epilogue end applies cannot be directly determined from the line number information alone; the function must be determined in combination with the subroutine information entries of the compilation (including inlined subroutines).*

*In the case of a trivial function, both prologue end and epilogue begin may occur at the same address.*

12. **DW\_LNS\_set\_isa**

The DW\_LNS\_set\_isa opcode takes a single unsigned LEB128 operand and stores that value in the isa register of the state machine.

13. **DW\_LNS\_extended\_op**

The DW\_LNS\_extended\_op opcode takes two operands. The first is an unsigned LEB128 value that gives the size of the operand that follows. The second begins with an extended opcode which is followed by operands appropriate to that opcode.

**6.2.5.3 Extended Opcodes**

Extended opcodes are used as part of a [DW\\_LNS\\_extended\\_op](#) operation (see Section [6.2.3 on page 158](#)).

The extended opcodes are as follows:

1. **DW\_LNE\_end\_sequence**

The DW\_LNE\_end\_sequence opcode takes no operands. It sets the end\_sequence register of the state machine to “true” and appends a row to the matrix using the current values of the state-machine registers. Then it resets the registers to the initial values specified above (see Section [6.2.2 on page 157](#)). Every line number program sequence must end with a [DW\\_LNE\\_end\\_sequence](#) instruction which creates a row whose address is that of the byte after the last target machine instruction of the sequence.

1     2. **DW\_LNE\_set\_address**

2     The DW\_LNE\_set\_address opcode takes a single relocatable address as an  
3     operand. The size of the operand is the size of an address on the target  
4     machine. It sets the address register to the value given by the relocatable  
5     address and sets the op\_index register to 0.

6     If the address value is the reserved target address (see Section 2.4.1 on  
7     page 26), no instructions are associated with subsequent rows up to but not  
8     including the subsequent DW\_LNE\_set\_address or DW\_LNE\_end\_sequence  
9     opcode, which is equivalent to omitting that sequence of opcodes.

10     *All of the other line number program opcodes that affect the address register add a  
11     delta to it. This instruction stores a relocatable value into the address register instead.*

12     3. **DW\_LNE\_set\_discriminator**

13     The DW\_LNE\_set\_discriminator opcode takes a single parameter, an  
14     unsigned LEB128 integer. It sets the discriminator register to the new value.

15     4. **DW\_LNE\_padding**

16     The DW\_LNE\_padding opcode is followed by a single operand which  
17     consists of a sequence of zero or more arbitrary bytes up to the length  
18     specified by the unsigned LEB128 integer that precedes all extended opcodes.  
19     The opcode and operand have no effect on the line number program.

20     *This permits a producer to pad or overwrite arbitrary parts of a line number program,  
21     with a minimum of the three bytes needed to encode any extended opcode.*

22     *Appendix D.5.3 on page 341 gives some sample line number programs.*

23     

## 6.3 Macro Information

24     *Some languages, such as C and C++, provide a way to replace text in the source program  
25     with macros defined either in the source file itself, or in another file included by the source  
26     file. Because these macros are not themselves defined in the target language, it is difficult  
27     to represent their definitions using the standard language constructs of DWARF. The  
28     debugging information therefore reflects the state of the source after the macro definition  
29     has been expanded, rather than as the programmer wrote it. The macro information table  
30     provides a way of preserving the original source in the debugging information.*

31     As described in Section 3.1.1 on page 62, the macro information for a given  
32     compilation unit is represented in the .debug\_macro section of an object file.

1 The macro information for each compilation unit consists of one or more macro  
2 units. Each macro unit starts with a header and is followed by a series of macro  
3 information entries or file inclusion entries. Each entry consists of an opcode  
4 followed by zero or more operands. Each macro unit ends with an entry  
5 containing an opcode of 0.

6 In all macro information entries, the line number of the entry is encoded as an  
7 unsigned LEB128 integer.

### 8 **6.3.1 Macro Information Header**

9 The macro information header contains the following fields:

10 1. `version` (uhalf)

11 A version number (see Section 7.23 on page 249). This number is specific to  
12 the macro information and is independent of the DWARF version number.

13 2. `flags` (ubyte)

14 The bits of the `flags` field are interpreted as a set of flags, some of which may  
15 indicate that additional fields follow.

16 The following flags, beginning with the least significant bit, are defined:

17 • `offset_size_flag`

18 If the `offset_size_flag` is zero, the header is for a 32-bit DWARF format  
19 macro section and all offsets are 4 bytes long; if it is one, the header is for  
20 a 64-bit DWARF format macro section and all offsets are 8 bytes long.

21 This flag does not apply to the the following entries:

22 [DW\\_MACRO\\_define\\_sup4](#), [DW\\_MACRO\\_define\\_sup8](#),  
23 [DW\\_MACRO\\_undef\\_sup4](#), [DW\\_MACRO\\_undef\\_sup8](#),  
24 [DW\\_MACRO\\_import\\_sup4](#) and [DW\\_MACRO\\_import\\_sup8](#).

25 • `debug_line_offset_flag`

26 If the `debug_line_offset_flag` is one, the `debug_line_offset` field (see  
27 below) is present. If zero, that field is omitted.

28 • `opcode_operands_table_flag`

29 If the `opcode_operands_table_flag` is one, the `opcode_operands_table`  
30 field (see below) is present. If zero, that field is omitted.

31 All other flags are reserved by DWARF.

## Chapter 6. Other Debugging Information

### 3. `debug_line_offset`

An offset in the `.debug_line` section (if this header is in a `.debug_macro` section) or `.debug_line.dwo` section (if this header is in a `.debug_macro.dwo` section) of the beginning of the line number information in the containing compilation, encoded as a 4-byte offset for a 32-bit DWARF format macro section and an 8-byte offset for a 64-bit DWARF format macro section.

### 4. `opcode_operands_table`

An `opcode_operands_table` describing the operands of the macro information entry opcodes.

The macro information entries defined in this standard may, but need not, be described in the table, while other user-defined entry opcodes used in the section are described there. Vendor extension entry opcodes are allocated in the range from `DW_MACRO_lo_user` to `DW_MACRO_hi_user`. Other unassigned codes are reserved for future DWARF standards.

The table starts with a 1-byte count of the defined opcodes, followed by an entry for each of those opcodes. Each entry starts with a 1-byte unsigned opcode number, followed by unsigned LEB128 encoded number of operands and for each operand there is a single unsigned byte describing the form in which the operand is encoded. The allowed forms are: `DW_FORM_block`, `DW_FORM_block1`, `DW_FORM_block2`, `DW_FORM_block4`, `DW_FORM_data1`, `DW_FORM_data2`, `DW_FORM_data4`, `DW_FORM_data8`, `DW_FORM_data16`, `DW_FORM_flag`, `DW_FORM_line_strp`, `DW_FORM_sdata`, `DW_FORM_sec_offset`, `DW_FORM_string`, `DW_FORM_strp`, `DW_FORM_strp8`, `DW_FORM_strp_sup`, `DW_FORM_strp_sup8`, `DW_FORM_strx`, `DW_FORM_strx1`, `DW_FORM_strx2`, `DW_FORM_strx3`, `DW_FORM_strx4` and `DW_FORM_udata`.

### 6.3.2 Macro Information Entries

All macro information entries within a `.debug_macro` section for a given compilation unit appear in the same order in which the directives were processed by the compiler (after taking into account the effect of the macro import directives).

*The source file in which a macro information entry occurs can be derived by interpreting the sequence of entries from the beginning of the `.debug_macro` section. `DW_MACRO_start_file` and `DW_MACRO_end_file` indicate changes in the containing file.*

1 **6.3.2.1 Define and Undefine Entries**

2 The define and undefine macro entries have multiple forms that use different  
3 representations of their two operands.

4 While described in pairs below, the forms of define and undefine entries may be  
5 freely intermixed.

6 1. **DW\_MACRO\_define, DW\_MACRO\_undef**

7 A DW\_MACRO\_define or DW\_MACRO\_undef entry has two operands. The  
8 first operand encodes the source line number of the #define or #undef macro  
9 directive. The second operand is a null-terminated character string for the  
10 macro being defined or undefined.

11 The contents of the operands are described below (see Sections 6.3.2.2 and  
12 6.3.2.3 following).

13 2. **DW\_MACRO\_define\_strp, DW\_MACRO\_undef\_strp**

14 A DW\_MACRO\_define\_strp or DW\_MACRO\_undef\_strp entry has two  
15 operands. The first operand encodes the source line number of the #define or  
16 #undef macro directive. The second operand consists of an offset into a string  
17 table contained in the .debug\_str section of the object file. The size of the  
18 operand is given in the header offset\_size\_flag field.

19 The contents of the operands are described below (see Sections 6.3.2.2 and  
20 6.3.2.3 following).

21 3. **DW\_MACRO\_define\_strx, DW\_MACRO\_undef\_strx**

22 A DW\_MACRO\_define\_strx or DW\_MACRO\_undef\_strx entry has two  
23 operands. The first operand encodes the line number of the #define or  
24 #undef macro directive. The second operand identifies a string; it is  
25 represented using an unsigned LEB128 encoded value, which is interpreted as  
26 a zero-based index into an array of offsets in the .debug\_str\_offsets section.

27 The contents of the operands are described below (see Sections 6.3.2.2 and  
28 6.3.2.3 following).

29 4. **DW\_MACRO\_define\_sup4, DW\_MACRO\_define\_sup8,  
30 DW\_MACRO\_undef\_sup4, DW\_MACRO\_undef\_sup8**

31 A DW\_MACRO\_define\_sup4, DW\_MACRO\_define\_sup8,  
32 DW\_MACRO\_undef\_sup4 or DW\_MACRO\_undef\_sup8 entry has two  
33 operands. The first operand encodes the line number of the #define or  
34 #undef macro directive. The second operand identifies a string; it is  
35 represented as an offset into a string table contained in the .debug\_str  
36 section of the supplementary object file. The size of the operand is 4-bytes for

1 DW\_MACRO\_define\_sup4 and DW\_MACRO\_undef\_sup4, and 8-bytes for  
2 DW\_MACRO\_define\_sup8 and DW\_MACRO\_undef\_sup8.

3 The contents of the operands are described below (see Sections 6.3.2.2 and  
4 6.3.2.3 following).

### 5 6.3.2.2 Macro Define String

6 In the case of a DW\_MACRO\_define, DW\_MACRO\_define\_strp,  
7 DW\_MACRO\_define\_strx, DW\_MACRO\_define\_sup4 or  
8 DW\_MACRO\_define\_sup8 entry, the value of the second operand is the name of  
9 the macro symbol that is defined at the indicated source line, followed  
10 immediately by the macro formal parameter list including the surrounding  
11 parentheses (in the case of a function-like macro) followed by the definition  
12 string for the macro. If there is no formal parameter list, then the name of the  
13 defined macro is followed immediately by its definition string.

14 In the case of a function-like macro definition, no whitespace characters appear  
15 between the name of the defined macro and the following left parenthesis.  
16 Formal parameters are separated by a comma without any whitespace. Exactly  
17 one space character separates the right parenthesis that terminates the formal  
18 parameter list and the following definition string.

19 In the case of a “normal” (that is, non-function-like) macro definition, exactly one  
20 space character separates the name of the defined macro from the following  
21 definition text.

### 22 6.3.2.3 Macro Undefine String

23 In the case of a DW\_MACRO\_undef, DW\_MACRO\_undef\_strp,  
24 DW\_MACRO\_undef\_strx, DW\_MACRO\_undef\_sup4 or  
25 DW\_MACRO\_undef\_sup8 entry, the value of the second string is the name of the  
26 pre-processor symbol that is undefined at the indicated source line.

### 27 6.3.2.4 Entries for Command Line Options

28 A DWARF producer generates a define or undefine entry for each pre-processor  
29 symbol which is defined or undefined by some means other than such a directive  
30 within the compiled source text. In particular, pre-processor symbol definitions  
31 and undefinitions which occur as a result of command line options (when  
32 invoking the compiler) are represented by their own define and undefine entries.

1 All such define and undefine entries representing compilation options appear  
2 before the first [DW\\_MACRO\\_start\\_file](#) entry for that compilation unit (see  
3 Section 6.3.3 following) and encode the value 0 in their line number operands.

### 4 **6.3.3 File Inclusion Entries**

#### 5 **6.3.3.1 Source Include Directives**

6 The following directives describe a source file inclusion directive (`#include` in  
7 C/C++) and the ending of an included file.

##### 8 **1. DW\_MACRO\_start\_file**

9 A `DW_MACRO_start_file` entry has two operands. The first operand encodes  
10 the line number of the source line on which the `#include` macro directive  
11 occurs. The second operand encodes a source file name index.

12 The source file name index is the file number in the line number information  
13 table for the compilation unit.

14 If a `DW_MACRO_start_file` entry is present, the header contains a reference  
15 to the `.debug_line` section or `.debug_line.dwo` section of the compilation, as  
16 appropriate.

##### 17 **2. DW\_MACRO\_end\_file**

18 A `DW_MACRO_end_file` entry has no operands. The presence of the entry  
19 marks the end of the current source file inclusion.

20 When providing macro information in an object file, a producer generates  
21 [DW\\_MACRO\\_start\\_file](#) and [DW\\_MACRO\\_end\\_file](#) entries for the source file  
22 submitted to the compiler for compilation. This [DW\\_MACRO\\_start\\_file](#) entry  
23 has the value 0 in its line number operand and references the file entry in the line  
24 number information table for the primary source file.

#### 25 **6.3.3.2 Importation of Macro Units**

26 The import entries make it possible to replicate macro units. The first form  
27 supports replication within the current compilation and the second form  
28 supports replication across separate executable or shared object files.

29 *Import entries do not reflect the source program and, in fact, are not necessary at all.*  
30 *However, they do provide a mechanism that can be used to reduce redundancy in the*  
31 *macro information and thereby to save space.*



1     **1. DW\_MACRO\_import**

2     A DW\_MACRO\_import entry has one operand, an offset into another part of  
3     the .debug\_macro section that is the beginning of a target macro unit. The size  
4     of the operand depends on the header offset\_size\_flag field. The  
5     DW\_MACRO\_import entry instructs the consumer to replicate the sequence  
6     of entries following the target macro header which begins at the given  
7     .debug\_macro offset, up to, but excluding, the terminating entry with opcode  
8     0, as though the sequence of entries occurs in place of the import operation.

9     **2. DW\_MACRO\_import\_sup4, DW\_MACRO\_import\_sup8**

10    A DW\_MACRO\_import\_sup4 or DW\_MACRO\_import\_sup8 entry has one  
11    operand, an offset from the start of the .debug\_macro section in the  
12    supplementary object file. The size of the operand is 4 bytes for  
13    DW\_MACRO\_import\_sup4 and 8 bytes for DW\_MACRO\_import\_sup8.  
14    Apart from the different location in which to find the macro unit, this entry  
15    type is equivalent to [DW\\_MACRO\\_import](#).

16    *These entry types are aimed at sharing duplicate macro units between*  
17    *.debug\_macro sections from different executable or shared object files.*

18    From within the .debug\_macro section of the supplementary object file,  
19    [DW\\_MACRO\\_define\\_strp](#) and [DW\\_MACRO\\_undef\\_strp](#) entries refer to the  
20    .debug\_str section of that same supplementary file; similarly,  
21    [DW\\_MACRO\\_import](#) entries refer to the .debug\_macro section of that same  
22    supplementary file.

23    **6.4 Call Frame Information**

24    *Debuggers often need to be able to view and modify the state of any subroutine activation*  
25    *that is on the call stack. An activation consists of:*

- 26    • *A code location that is within the subroutine. This location is either the place where*  
27    *the program stopped when the debugger got control (for example, a breakpoint), or*  
28    *is a place where a subroutine made a call or was interrupted by an asynchronous*  
29    *event (for example, a signal).*
- 30    • *An area of memory that is allocated on a stack called a “call frame.” The call frame*  
31    *is identified by an address on the stack. We refer to this address as the Canonical*  
32    *Frame Address or CFA. Typically, the CFA is defined to be the value of the stack*  
33    *pointer at the call site in the previous frame (which may be different from its value*  
34    *on entry to the current frame).*
- 35    • *A set of registers that are in use by the subroutine at the code location.*



## Chapter 6. Other Debugging Information

1 Typically, a set of registers are designated to be preserved across a call. If a callee wishes  
2 to use such a register, it saves the value that the register had at entry time in its call frame  
3 and restores it on exit. The code that allocates space on the call frame stack and performs  
4 the save operation is called the subroutine's prologue, and the code that performs the  
5 restore operation and deallocates the frame is called its epilogue. Typically, the prologue  
6 code is physically at the beginning of a subroutine and the epilogue code is at the end.

7 To be able to view or modify an activation that is not on the top of the call frame stack, the  
8 debugger must virtually unwind the stack of activations until it finds the activation of  
9 interest. A debugger virtually unwinds a stack in steps. Starting with the current  
10 activation it virtually restores any registers that were preserved by the current activation  
11 and computes the predecessor's CFA and code location. This has the logical effect of  
12 returning from the current subroutine to its predecessor. We say that the debugger  
13 virtually unwinds the stack because the actual state of the target process is unchanged.

14 The virtual unwind operation needs to know where registers are saved and how to  
15 compute the predecessor's CFA and code location. When considering an  
16 architecture-independent way of encoding this information one has to consider a number  
17 of special things:

- 18 • Prologue and epilogue code is not always in distinct blocks at the beginning and  
19 end of a subroutine. It is common to duplicate the epilogue code at the site of each  
20 return from the code. Sometimes a compiler breaks up the register save/unsave  
21 operations and moves them into the body of the subroutine to just where they are  
22 needed.
- 23 • Compilers use different ways to manage the call frame. Sometimes they use a frame  
24 pointer register, sometimes not.
- 25 • The algorithm to compute CFA changes as you progress through the prologue and  
26 epilogue code. (By definition, the CFA value does not change.)
- 27 • Some subroutines have no call frame.
- 28 • Sometimes a register is saved in another register that by convention does not need  
29 to be saved.
- 30 • Some architectures have special instructions that perform some or all of the register  
31 management in one instruction, leaving special information on the stack that  
32 indicates how registers are saved.
- 33 • Some architectures treat return address values specially. For example, in one  
34 architecture, the call instruction guarantees that the low order two bits will be zero  
35 and the return instruction ignores those bits. This leaves two bits of storage that  
36 are available to other uses that must be treated specially.

### 6.4.1 Structure of Call Frame Information

DWARF supports virtual unwinding by defining an architecture independent basis for recording how subprograms save and restore registers during their lifetimes. This basis must be augmented on some machines with specific information that is defined by an architecture specific ABI authoring committee, a hardware vendor, or a compiler producer. The body defining a specific augmentation is referred to below as the “augmenter.”

Abstractly, this mechanism describes a very large table that has the following structure:

```

LOC CFA R0 R1 ... RN
LO
L1
...
LN
```

The first column indicates an address for every location that contains code in a program. (In shared object files, this is an object-relative offset.) The remaining columns contain virtual unwinding rules that are associated with the indicated location.

The CFA column defines the rule which computes the Canonical Frame Address value; the rule may indicate either a register and a signed offset that are added together, or a DWARF expression that is evaluated.

The remaining columns are labelled by register number. This includes some registers that have special designation on some architectures such as the PC and the stack pointer register. (The actual mapping of registers for a particular architecture is defined by the augmenter.) The register columns contain rules that describe whether a given register has been saved and the rule to find the value for the register in the previous frame.

## Chapter 6. Other Debugging Information

1 The register rules are:

undefined	A register that has this rule has no recoverable value in the previous frame. (By convention, it is not preserved by a callee.)
same value	This register has not been modified from the previous frame. (By convention, it is preserved by the callee, but the callee has not modified it.)
offset(N)	The previous value of this register is saved at the address CFA+N where CFA is the current CFA value and N is a signed offset.
val_offset(N)	The previous value of this register is the value CFA+N where CFA is the current CFA value and N is a signed offset.
register(R)	The previous value of this register is stored in another register numbered R.
expression(E)	The previous value of this register is located at the address produced by executing the DWARF expression E (see <a href="#">Section 2.5 on page 26</a> ).
val_expression(E)	The previous value of this register is the value produced by executing the DWARF expression E (see <a href="#">Section 2.5 on page 26</a> ).
architectural	The rule is defined externally to this specification by the augments.

2 *This table would be extremely large if actually constructed as described. Most of the*  
3 *entries at any point in the table are identical to the ones above them. The whole table can*  
4 *be represented quite compactly by recording just the differences starting at the beginning*  
5 *address of each subroutine in the program.*

6 The virtual unwind information is encoded in a self-contained section called  
7 `.debug_frame`. Entries in a `.debug_frame` section are aligned on a multiple of the  
8 address size relative to the start of the section and come in two forms: a Common  
9 Information Entry (CIE) and a Frame Description Entry (FDE).

10 *If the range of code addresses for a function is not contiguous, there may be multiple CIEs*  
11 *and FDEs corresponding to the parts of that function.*

## Chapter 6. Other Debugging Information

1 A Common Information Entry holds information that is shared among many  
2 Frame Description Entries. There is at least one CIE in every non-empty  
3 `.debug_frame` section. A CIE contains the following fields, in order:

4 1. `length` (initial length)

5 A constant that gives the number of bytes of the CIE structure, not including  
6 the length field itself (see Section 7.2.2 on page 192). The size of the length  
7 field plus the value of `length` must be an integral multiple of the address size.

8 2. `CIE_id` (4 or 8 bytes, see Section 7.4 on page 204)

9 A constant that is used to distinguish CIEs from FDEs.

10 3. `version` (ubyte)

11 A version number (see Section 7.24 on page 251). This number is specific to  
12 the call frame information and is independent of the DWARF version number.

13 4. `augmentation` (sequence of UTF-8 characters)

14 A null-terminated UTF-8 string that identifies the augmentation to this CIE or  
15 to the FDEs that use it. If a reader encounters an augmentation string that is  
16 unexpected, then only the following fields can be read:

- 17 • CIE: `length`, `CIE_id`, `version`, `augmentation`
- 18 • FDE: `length`, `CIE_pointer`, `initial_location`, `address_range`

19 If there is no augmentation, this value is a zero byte.

20 *The augmentation string allows users to indicate that there is additional*  
21 *target-specific information in the CIE or FDE which is needed to virtually unwind a*  
22 *stack frame. For example, this might be information about dynamically allocated data*  
23 *which needs to be freed on exit from the routine.*

24 *Because the `.debug_frame` section is useful independently of any `.debug_info`*  
25 *section, the augmentation string always uses UTF-8 encoding.*

26 5. `address_size` (ubyte)

27 The size of a target address in bytes in this CIE and any FDEs that use it. If a  
28 compilation unit exists for this frame, its address size must match the address  
29 size here.

30 6. `reserved`<sup>3</sup> (ubyte, MBZ)

31 7. `code_alignment_factor` (unsigned LEB128)

32 A constant that is factored out of all advance location instructions (see Section  
33 6.4.2.1 on page 184). The resulting value is *operand* \* `code_alignment_factor`.

---

<sup>3</sup>This allows backward compatible support of the deprecated `segment_selector_size` field which was defined in DWARF Version 5 and earlier.

## Chapter 6. Other Debugging Information

- 1 8. `data_alignment_factor` (signed LEB128)  
2 A constant that is factored out of certain offset instructions (see Sections  
3 [6.4.2.2 on page 185](#) and [6.4.2.3 on page 186](#)). The resulting value is  
4  $operand * data\_alignment\_factor$ .
  - 5 9. `return_address_register` (unsigned LEB128)  
6 An unsigned LEB128 constant that indicates which column in the rule table  
7 represents the return address of the function. Note that this column might not  
8 correspond to an actual machine register.
  - 9 10. `initial_instructions` (array of ubyte)  
10 A sequence of rules that are interpreted to create the initial setting of each  
11 column in the table.  
12 The default rule for all columns before interpretation of the initial instructions  
13 is the undefined rule. However, an ABI authoring body or a compilation  
14 system authoring body may specify an alternate default value for any or all  
15 columns.
  - 16 11. `padding` (array of ubyte)  
17 Enough `DW_CFA_nop` instructions to make the size of this entry match the  
18 length value above.
- 19 An FDE contains the following fields, in order:
- 20 1. `length` ([initial length](#))  
21 A constant that gives the number of bytes of the header and instruction  
22 stream for this function, not including the length field itself (see Section [7.2.2](#)  
23 [on page 192](#)). The size of the length field plus the value of length must be an  
24 integral multiple of the address size.
  - 25 2. `CIE_pointer` (4 or 8 bytes, see Section [7.4 on page 204](#))  
26 A constant offset into the `.debug_frame` section that denotes the CIE that is  
27 associated with this FDE.
  - 28 3. `initial_location` (target address)  
29 The address of the first location associated with this table entry.
  - 30 4. `address_range` (target address)  
31 The number of bytes of program instructions described by this entry.
  - 32 5. `instructions` (array of ubyte)  
33 A sequence of table defining instructions that are described in Section [6.4.2 on](#)  
34 [the next page](#).

- 1 6. padding (array of ubyte)
- 2 Enough `DW_CFA_nop` instructions to make the size of this entry match the
- 3 length value above.

## 4 6.4.2 Call Frame Instructions

5 Each call frame instruction is defined to take 0 or more operands. Some of the  
6 operands may be encoded as part of the opcode (see Section 7.24 on page 251).  
7 The instructions are defined in the following sections.

8 Some call frame instructions have operands that are encoded as DWARF  
9 expressions (see Section 2.5.1 on page 27). The following DWARF operators  
10 cannot be used in such operands:

- 11 • `DW_OP_addrx`, `DW_OP_call2`, `DW_OP_call4`, `DW_OP_call_ref`,  
12 `DW_OP_const_type`, `DW_OP_constx`, `DW_OP_convert`,  
13 `DW_OP_deref_type`, `DW_OP_regval_type` and `DW_OP_reinterpret`  
14 operators are not allowed in an operand of these instructions because the  
15 call frame information must not depend on other debug sections.
- 16 • `DW_OP_push_object_address` is not meaningful in an operand of these  
17 instructions because there is no object context to provide a value to push.
- 18 • `DW_OP_call_frame_cfa` is not meaningful in an operand of these  
19 instructions because its use would be circular.

20 *Call frame instructions to which these restrictions apply include*  
21 *`DW_CFA_def_cfa_expression`, `DW_CFA_expression` and `DW_CFA_val_expression`.*

### 22 6.4.2.1 Row Creation Instructions

#### 23 1. `DW_CFA_set_loc`

24 The `DW_CFA_set_loc` instruction takes a single operand that represents a  
25 target address. The required action is to create a new table row using the  
26 specified address as the location. All other values in the new row are initially  
27 identical to the current row. The new location value is always greater than the  
28 current one.

#### 29 2. `DW_CFA_advance_loc`

30 The `DW_CFA_advance_loc` instruction takes a single operand (encoded with  
31 the opcode) that represents a constant delta. The required action is to create a  
32 new table row with a location value that is computed by taking the current  
33 entry's location value and adding the value of  $\text{delta} * \text{code\_alignment\_factor}$ .  
34 All other values in the new row are initially identical to the current row

1     3. **DW\_CFA\_advance\_loc1**

2     The DW\_CFA\_advance\_loc1 instruction takes a single ubyte operand that  
3     represents a constant delta. This instruction is identical to  
4     DW\_CFA\_advance\_loc except for the encoding and size of the delta operand.

5     4. **DW\_CFA\_advance\_loc2**

6     The DW\_CFA\_advance\_loc2 instruction takes a single uhalf operand that  
7     represents a constant delta. This instruction is identical to  
8     DW\_CFA\_advance\_loc except for the encoding and size of the delta operand.

9     5. **DW\_CFA\_advance\_loc4**

10    The DW\_CFA\_advance\_loc4 instruction takes a single uword operand that  
11    represents a constant delta. This instruction is identical to  
12    DW\_CFA\_advance\_loc except for the encoding and size of the delta operand.

13    **6.4.2.2 CFA Definition Instructions**

14    1. **DW\_CFA\_def\_cfa**

15    The DW\_CFA\_def\_cfa instruction takes two unsigned LEB128 operands  
16    representing a register number and a (non-factored) offset. The required  
17    action is to define the current CFA rule to use the provided register and offset.

18    2. **DW\_CFA\_def\_cfa\_sf**

19    The DW\_CFA\_def\_cfa\_sf instruction takes two operands: an unsigned  
20    LEB128 value representing a register number and a signed LEB128 factored  
21    offset. This instruction is identical to DW\_CFA\_def\_cfa except that the second  
22    operand is signed and factored. The resulting offset is  
23    *factored\_offset* \* *data\_alignment\_factor*.

24    3. **DW\_CFA\_def\_cfa\_register**

25    The DW\_CFA\_def\_cfa\_register instruction takes a single unsigned LEB128  
26    operand representing a register number. The required action is to define the  
27    current CFA rule to use the provided register (but to keep the old offset). This  
28    operation is valid only if the current CFA rule is defined to use a register and  
29    offset.

30    4. **DW\_CFA\_def\_cfa\_offset**

31    The DW\_CFA\_def\_cfa\_offset instruction takes a single unsigned LEB128  
32    operand representing a (non-factored) offset. The required action is to define  
33    the current CFA rule to use the provided offset (but to keep the old register).  
34    This operation is valid only if the current CFA rule is defined to use a register  
35    and offset.

1       5. **DW\_CFA\_def\_cfa\_offset\_sf**

2       The DW\_CFA\_def\_cfa\_offset\_sf instruction takes a signed LEB128 operand  
3       representing a factored offset. This instruction is identical to  
4       [DW\\_CFA\\_def\\_cfa\\_offset](#) except that the operand is signed and factored. The  
5       resulting offset is *factored\_offset* \* *data\_alignment\_factor*. This operation is  
6       valid only if the current CFA rule is defined to use a register and offset.

7       6. **DW\_CFA\_def\_cfa\_expression**

8       The DW\_CFA\_def\_cfa\_expression instruction takes a single operand encoded  
9       as an [exprval](#) value representing a DWARF expression. The required action is  
10      to establish that expression as the means by which the current CFA is  
11      computed.

12      *See Section 6.4.2 on page 184 regarding restrictions on the DWARF expression*  
13      *operators that can be used.*

14      **6.4.2.3 Register Rule Instructions**

15      1. **DW\_CFA\_undefined**

16      The DW\_CFA\_undefined instruction takes a single unsigned LEB128 operand  
17      that represents a register number. The required action is to set the rule for the  
18      specified register to “undefined.”

19      2. **DW\_CFA\_same\_value**

20      The DW\_CFA\_same\_value instruction takes a single unsigned LEB128  
21      operand that represents a register number. The required action is to set the  
22      rule for the specified register to “same value.”

23      3. **DW\_CFA\_offset**

24      The DW\_CFA\_offset instruction takes two operands: a register number  
25      (encoded with the opcode) and an unsigned LEB128 constant representing a  
26      factored offset. The required action is to change the rule for the register  
27      indicated by the register number to be an offset(N) rule where the value of N  
28      is *factored\_offset* \* *data\_alignment\_factor*.

29      4. **DW\_CFA\_offset\_extended**

30      The DW\_CFA\_offset\_extended instruction takes two unsigned LEB128  
31      operands representing a register number and a factored offset. This  
32      instruction is identical to [DW\\_CFA\\_offset](#) except for the encoding and size of  
33      the register operand.



1     5. **DW\_CFA\_offset\_extended\_sf**

2     The DW\_CFA\_offset\_extended\_sf instruction takes two operands: an  
3     unsigned LEB128 value representing a register number and a signed LEB128  
4     factored offset. This instruction is identical to [DW\\_CFA\\_offset\\_extended](#)  
5     except that the second operand is signed and factored. The resulting offset is  
6     *factored\_offset* \* *data\_alignment\_factor*.

7     6. **DW\_CFA\_val\_offset**

8     The DW\_CFA\_val\_offset instruction takes two unsigned LEB128 operands  
9     representing a register number and a factored offset. The required action is to  
10    change the rule for the register indicated by the register number to be a  
11    *val\_offset(N)* rule where the value of N is  
12    *factored\_offset* \* *data\_alignment\_factor*.

13    7. **DW\_CFA\_val\_offset\_sf**

14    The DW\_CFA\_val\_offset\_sf instruction takes two operands: an unsigned  
15    LEB128 value representing a register number and a signed LEB128 factored  
16    offset. This instruction is identical to [DW\\_CFA\\_val\\_offset](#) except that the  
17    second operand is signed and factored. The resulting offset is  
18    *factored\_offset* \* *data\_alignment\_factor*.

19    8. **DW\_CFA\_register**

20    The DW\_CFA\_register instruction takes two unsigned LEB128 operands  
21    representing register numbers. The required action is to set the rule for the  
22    first register to be *register(R)* where R is the second register.

23    9. **DW\_CFA\_expression**

24    The DW\_CFA\_expression instruction takes two operands: an unsigned  
25    LEB128 value representing a register number, and an [exprval](#) value  
26    representing a DWARF expression. The required action is to change the rule  
27    for the register indicated by the register number to be an *expression(E)* rule  
28    where E is the DWARF expression. That is, the DWARF expression computes  
29    the address. The value of the CFA is pushed on the DWARF evaluation stack  
30    prior to execution of the DWARF expression.

31    *See Section 6.4.2 on page 184 regarding restrictions on the DWARF expression*  
32    *operators that can be used.*

10. **DW\_CFA\_val\_expression**

The DW\_CFA\_val\_expression instruction takes two operands: an unsigned LEB128 value representing a register number, and an `exprval` value representing a DWARF expression. The required action is to change the rule for the register indicated by the register number to be a `val_expression(E)` rule where E is the DWARF expression. That is, the DWARF expression computes the value of the given register. The value of the CFA is pushed on the DWARF evaluation stack prior to execution of the DWARF expression.

*See Section 6.4.2 on page 184 regarding restrictions on the DWARF expression operators that can be used.*

11. **DW\_CFA\_restore**

The DW\_CFA\_restore instruction takes a single operand (encoded with the opcode) that represents a register number. The required action is to change the rule for the indicated register to the rule assigned it by the `initial_instructions` in the CIE.

12. **DW\_CFA\_restore\_extended**

The DW\_CFA\_restore\_extended instruction takes a single unsigned LEB128 operand that represents a register number. This instruction is identical to `DW_CFA_restore` except for the encoding and size of the register operand.

**6.4.2.4 Row State Instructions**

*The next two instructions provide the ability to stack and retrieve complete register states. They may be useful, for example, for a compiler that moves epilogue code into the body of a function.*

1. **DW\_CFA\_remember\_state**

The DW\_CFA\_remember\_state instruction takes no operands. The required action is to push the set of rules for the current CFA and every register onto an implicit stack.

2. **DW\_CFA\_restore\_state**

The DW\_CFA\_restore\_state instruction takes no operands. The required action is to pop the set of rules off the implicit stack and place them in the current row.

**6.4.2.5 Padding Instruction**

1. **DW\_CFA\_nop**

The DW\_CFA\_nop instruction has no operands and no required actions. It is used as padding to make a CIE or FDE an appropriate size.

### 6.4.3 Call Frame Instruction Usage

To determine the virtual unwind rule set for a given location (L1), search through the FDE headers looking at the *initial\_location* and *address\_range* values to see if L1 is contained in the FDE. If so, then:

1. Initialize a register set by reading the *initial\_instructions* field of the associated CIE. Set L2 to the value of the *initial\_location* field from the FDE header.
2. Read and process the FDE's instruction sequence until a *DW\_CFA\_advance\_loc*, *DW\_CFA\_set\_loc*, or the end of the instruction stream is encountered.
3. If a *DW\_CFA\_advance\_loc* or *DW\_CFA\_set\_loc* instruction is encountered, then compute a new location value (L2). If  $L1 \geq L2$  then process the instruction and go back to step 2.
4. The end of the instruction stream can be thought of as a *DW\_CFA\_set\_loc* (*initial\_location* + *address\_range*) instruction. Note that the FDE is ill-formed if L2 is less than L1.

The rules in the register set now apply to location L1.

For an example, see Appendix D.6 on page 343.

### 6.4.4 Call Frame Calling Address

When virtually unwinding frames, consumers frequently wish to obtain the address of the instruction which called a subroutine. This information is not always provided. Typically, however, one of the registers in the virtual unwind table is the Return Address.

If a Return Address register is defined in the virtual unwind table, and its rule is undefined (for example, by *DW\_CFA\_undefined*), then there is no return address and no call address, and the virtual unwind of stack activations is complete.

In most cases the return address is in the same context as the calling address, but that need not be the case, especially if the producer knows in some way the call never will return. The context of the 'return address' might be on a different line, in a different lexical *block*, or past the end of the calling subroutine. If a consumer were to assume that it was in the same context as the calling address, the virtual unwind might fail.

For architectures with constant-length instructions where the return address immediately follows the call instruction, a simple solution is to subtract the length of an instruction from the return address to obtain the calling instruction. For architectures with variable-length instructions (for example, x86), this is not possible. However, subtracting 1 from the return address, although not guaranteed to provide the exact calling address, generally will produce an address within the same context as the calling address, and that usually is sufficient.

## Chapter 6. Other Debugging Information

*(empty page)*

# Chapter 7

## Data Representation

This section describes the binary representation of the debugging information entry itself, of the attribute types and of other fundamental elements described above.

### 7.1 Vendor Extensibility

To reserve a portion of the DWARF name space and ranges of enumeration values for use for vendor specific extensions, special labels are reserved for tag names, attribute names, base type encodings, location operations, language names, calling conventions and call frame instructions.

The labels denoting the beginning and end of the reserved value range for vendor specific extensions consist of the appropriate prefix (DW\_AT, DW\_ATE, DW\_CC, DW\_CFA, DW\_END, DW\_IDX, DW\_LLE, DW\_LNAME, DW\_LNCT, DW\_LNE, DW\_MACRO, DW\_OP, DW\_RLE, DW\_TAG, DW\_UT) followed by `_lo_user` or `_hi_user`. Values in the range between `prefix_lo_user` and `prefix_hi_user` inclusive, are reserved for vendor specific extensions. Vendors may use values in this range without conflicting with current or future system-defined values. All other values are reserved for use by the system.

*For example, for debugging information entry tags, the special labels are `DW_TAG_lo_user` and `DW_TAG_hi_user`.*

*There may also be codes for vendor specific extensions between the number of standard line number opcodes and the first special line number opcode. However, since the number of standard opcodes varies with the DWARF version, the range for extensions is also version dependent. Thus, `DW_LNS_lo_user` and `DW_LNS_hi_user` symbols are not defined.*

1 Vendor defined tags, attributes, base type encodings, location atoms, language  
2 names, line number actions, calling conventions and call frame instructions,  
3 conventionally use the form `prefix_vendor_id_name`, where *vendor\_id* is some  
4 identifying character sequence chosen so as to avoid conflicts with other vendors.

5 To ensure that extensions added by one vendor may be safely ignored by  
6 consumers that do not understand those extensions, the following rules must be  
7 followed:

- 8 1. New attributes are added in such a way that a debugger may recognize the  
9 format of a new attribute value without knowing the content of that attribute  
10 value.
- 11 2. The semantics of any new attributes do not alter the semantics of previously  
12 existing attributes.
- 13 3. The semantics of any new tags do not conflict with the semantics of  
14 previously existing tags.
- 15 4. New forms of attribute value are not added.

## 16 **7.2 Reserved Values**

### 17 **7.2.1 Error Values**

18 As a convenience for consumers of DWARF information, the value 0 is reserved  
19 in the encodings for attribute names, attribute forms, base type encodings,  
20 location operations, languages, line number program opcodes, macro  
21 information entries and tag names to represent an error condition or unknown  
22 value. DWARF does not specify names for these reserved values, because they  
23 do not represent valid encodings for the given type and do not appear in  
24 DWARF debugging information.

### 25 **7.2.2 Initial Length Values**

26 An initial length field is one of the fields that occur at the beginning of those  
27 DWARF sections that have a header (`.debug_aranges`, `.debug_info`,  
28 `.debug_line`, `.debug_loclists`, `.debug_names` and `.debug_rnglists`) or the  
29 length field that occurs at the beginning of the CIE and FDE structures in the  
30 `.debug_frame` section.

1 In an initial length field, the values 0xffffffff0 through 0xffffffff are reserved  
2 by DWARF to indicate some form of extension relative to DWARF Version 2;  
3 such values must not be interpreted as a length field. The use of one such value,  
4 0xffffffff, is defined in Section 7.4 on page 204; the use of the other values is  
5 reserved for possible future extensions.

## 6 **7.3 Relocatable, Split, Executable, Shared, Package** 7 **and Supplementary Object Files**

### 8 **7.3.1 Relocatable Object Files**

9 A DWARF producer (for example, a compiler) typically generates its debugging  
10 information as part of a relocatable object file. Relocatable object files are then  
11 combined by a linker to form an executable file. During the linking process, the  
12 linker resolves (binds) symbolic references between the various object files, and  
13 relocates the contents of each object file into a combined virtual address space.

14 The DWARF debugging information is placed in several sections (see Appendix  
15 [B on page 284](#)), and requires an object file format capable of representing these  
16 separate sections. There are symbolic references between these sections, and also  
17 between the debugging information sections and the other sections that contain  
18 the text and data of the program itself. Many of these references require  
19 relocation, and the producer must emit the relocation information appropriate to  
20 the object file format and the target processor architecture. These references  
21 include the following:

- 22 • The compilation unit header (see Section [7.5.1 on page 209](#)) in the  
23 `.debug_info` section contains a reference to the `.debug_abbrev` table. This  
24 reference requires a relocation so that after linking, it refers to that  
25 contribution to the combined `.debug_abbrev` section in the executable file.
- 26 • Debugging information entries may have attributes with the form  
27 `DW_FORM_addr` (see Section [7.5.4 on page 218](#)). These attributes represent  
28 locations within the virtual address space of the program, and require  
29 relocation.
- 30 • A DWARF expression may contain a `DW_OP_addr` (see Section [2.5.1.1 on](#)  
31 [page 27](#)) which contains a location within the virtual address space of the  
32 program, and require relocation.

## Chapter 7. Data Representation

- 1       • Debugging information entries may have attributes with the form  
2       [DW\\_FORM\\_sec\\_offset](#) (see Section 7.5.4 on page 218). These attributes refer  
3       to debugging information in other debugging information sections within  
4       the object file, and must be relocated during the linking process.
- 5       • Debugging information entries may have attributes with the form  
6       [DW\\_FORM\\_ref\\_addr](#) (see Section 7.5.4 on page 218). These attributes refer  
7       to debugging information entries that may be outside the current  
8       compilation unit. These values require both symbolic binding and  
9       relocation.
- 10       • Debugging information entries may have attributes with the form  
11       [DW\\_FORM\\_strp](#) or [DW\\_FORM\\_strp8](#) (see Section 7.5.4 on page 218).  
12       These attributes refer to strings in the `.debug_str` section. These values  
13       require relocation.
- 14       • The `.debug_macro` section may have [DW\\_MACRO\\_define\\_strp](#) and  
15       [DW\\_MACRO\\_undef\\_strp](#) entries (see Section 6.3.2.1 on page 175). These  
16       entries refer to strings in the `.debug_str` section. These values require  
17       relocation.
- 18       • Entries in the `.debug_addr` and `.debug_aranges` sections may contain  
19       references to locations within the virtual address space of the program, and  
20       thus require relocation.
- 21       • Entries in the `.debug_loclists` and `.debug_rnglists` sections may contain  
22       references to locations within the virtual address space of the program  
23       depending on whether certain kinds of location or range list entries are  
24       used, and thus require relocation.
- 25       • In the `.debug_line` section, the operand of the [DW\\_LNE\\_set\\_address](#)  
26       opcode is a reference to a location within the virtual address space of the  
27       program, and requires relocation.
- 28       • The `.debug_str_offsets` section contains a list of string offsets, each of  
29       which is an offset of a string in the `.debug_str` section. Each of these offsets  
30       requires relocation. Depending on the implementation, these relocations  
31       may be implicit (that is, the producer may not need to emit any explicit  
32       relocation information for these offsets).
- 33       • The `debug_info_offset` field in the `.debug_aranges` header and the list of  
34       compilation units following the `.debug_names` header contain references to  
35       the `.debug_info` section. These references require relocation so that after  
36       linking they refer to the correct contribution in the combined `.debug_info`  
37       section in the executable file.



- Frame descriptor entries in the `.debug_frame` section (see Section 6.4.1 on page 180) contain an `initial_location` field value within the virtual address space of the program and require relocation.

*Note that operands of classes `constant` and `flag` do not require relocation. Attribute operands that use forms `DW_FORM_string`, `DW_FORM_ref1`, `DW_FORM_ref2`, `DW_FORM_ref4`, `DW_FORM_ref8`, or `DW_FORM_ref_udata` also do not need relocation.*

## 7.3.2 Split DWARF Object Files

A DWARF producer may partition the debugging information such that the majority of the debugging information can remain in individual object files without being processed by the linker.

*This reduces link time by reducing the amount of information the linker must process.*

### 7.3.2.1 First Partition (with Skeleton Unit)

The first partition contains debugging information that must still be processed by the linker, and includes the following:

- The line number tables, frame tables, and accelerated access tables, in the usual sections: `.debug_line`, `.debug_line_str`, `.debug_frame`, `.debug_names` and `.debug_aranges`, respectively.
- An address table, in the `.debug_addr` section. This table contains all addresses and constants that require link-time relocation, and items in the table can be referenced indirectly from the debugging information via the `DW_FORM_addrx`, `DW_FORM_addrx1`, `DW_FORM_addrx2`, `DW_FORM_addrx3` and `DW_FORM_addrx4` forms, by the `DW_OP_addrx` and `DW_OP_constx` operators, and by certain of the `DW_LLE_*` location list and `DW_RLE_*` range list entries.
- A skeleton compilation unit, as described in Section 3.1.2 on page 69, in the `.debug_info` section.
- An abbreviations table for the skeleton compilation unit, in the `.debug_abbrev` section used by the `.debug_info` section.
- A string table, in the `.debug_str` section. The string table is necessary only if the skeleton compilation unit uses one of the indirect string forms (`DW_FORM_strp`, `DW_FORM_strp8`, `DW_FORM_strx`, `DW_FORM_strx1`, `DW_FORM_strx2`, `DW_FORM_strx3` or `DW_FORM_strx4`).

- A string offsets table, in the `.debug_str_offsets` section for strings in the `.debug_str` section. The string offsets table is necessary only if the skeleton compilation unit uses one of the indexed string forms (`DW_FORM_strx`, `DW_FORM_strx1`, `DW_FORM_strx2`, `DW_FORM_strx3`, `DW_FORM_strx4`).

The attributes contained in the skeleton compilation unit can be used by a DWARF consumer to find the DWARF object file that contains the second partition.

### 7.3.2.2 Second Partition (Unlinked or in a `.dwo` File)

The second partition contains the debugging information that does not need to be processed by the linker. These sections may be left in the object files and ignored by the linker (that is, not combined and copied to the executable object file), or they may be placed by the producer in a separate DWARF object file. This partition includes the following:

- The full compilation unit, in the `.debug_info.dwo` section.  
Attributes contained in the full compilation unit may refer to machine addresses indirectly using one of the `DW_FORM_addrx`, `DW_FORM_addrx1`, `DW_FORM_addrx2`, `DW_FORM_addrx3` or `DW_FORM_addrx4` forms, which access the table of addresses specified by the `DW_AT_addr_base` attribute in the associated skeleton unit. Location descriptions may similarly do so using the `DW_OP_addrx` and `DW_OP_constx` operations.
- Separate type units, in the `.debug_info.dwo` section.
- Abbreviations table(s) for the compilation unit and type units, in the `.debug_abbrev.dwo` section used by the `.debug_info.dwo` section.
- Location lists, in the `.debug_loclists.dwo` section.
- Range lists, in the `.debug_rnglists.dwo` section.
- A specialized line number table (for the type units, and macro information), in the `.debug_line.dwo` section.

This table contains only the directory and filename lists needed to interpret `DW_AT_decl_file` attributes in the debugging information entries and `DW_MACRO_start_file` entries in the macro information.

- Macro information, in the `.debug_macro.dwo` section.
- A string table, in the `.debug_str.dwo` section.

- A string offsets table, in the `.debug_str_offsets.dwo` section for the strings in the `.debug_str.dwo` section.

Attributes that refer to the `.debug_str.dwo` string table do so only indirectly through the `.debug_str_offsets.dwo` section using the forms `DW_FORM_strx`, `DW_FORM_strx1`, `DW_FORM_strx2`, `DW_FORM_strx3` or `DW_FORM_strx4`, or the macro entries `DW_MACRO_define_strx` or `DW_MACRO_undef_strx`. Direct reference (for example, using forms `DW_FORM_strp` or `DW_FORM_strp8`, or the macro entries `DW_MACRO_define_strp` or `DW_MACRO_undef_strp`) is not allowed.

Except where noted otherwise, all references in this document to a debugging information section (for example, `.debug_info`), apply also to the corresponding split DWARF section (for example, `.debug_info.dwo`).

Split DWARF object files do not get linked with any other files, therefore references between sections must not make use of normal object file relocation information. As a result, symbolic references within or between sections (such as from using `DW_FORM_ref_addr` and `DW_OP_call_ref`) are not possible. Split DWARF object files contain at most one compilation unit.

### 7.3.3 Executable Objects and `.dwo` Files

The relocated addresses in the debugging information for an executable object are virtual addresses.

The sections containing the debugging information are typically not loaded as part of the memory image of the program (in ELF terminology, the sections are not "allocatable" and are not part of a loadable segment). Therefore, the debugging information sections described in this document are typically linked as if they were each to be loaded at virtual address 0. Similarly, debugging information in a `.dwo` file is not loaded in the memory image. The absence (or non-use) of relocation information in a `.dwo` file means that sections described in this document are effectively linked as if they were each to be loaded at virtual address 0. In both cases, references within the debugging information always implicitly indicate which section a particular offset refers to. (For example, a reference of form `DW_FORM_sec_offset` may refer to one of several sections, depending on the class allowed by a particular attribute of a debugging information entry, as shown in Table 7.5 on page 218.)

### 7.3.4 Shared Object Files

The relocated addresses in the debugging information for a shared object file are offsets relative to the start of the lowest region of memory loaded from that shared object file.

*This requirement makes the debugging information for shared object files position independent. Virtual addresses in a shared object file may be calculated by adding the offset to the base address at which the object file was attached. This offset is available in the run-time linker's data structures.*

As with executable objects, the sections containing debugging information are typically not loaded as part of the memory image of the shared object, and are typically linked as if they were each to be loaded at virtual address 0.

### 7.3.5 DWARF Package Files

*Using split DWARF object files allows the developer to compile, link, and debug an application quickly with less link-time overhead, but a more convenient format is needed for saving the debug information for later debugging of a deployed application. A DWARF package file can be used to collect the debugging information from the object (or separate DWARF object) files produced during the compilation of an application.*

*The package file is typically placed in the same directory as the application, and is given the same name with a ".dwp" extension.*

A DWARF package file is itself an object file, using the same object file format (including byte order) as the corresponding application binary. It contains a file header, a section table, a number of DWARF debug information sections, and two index sections.

Each DWARF package file contains no more than one of each of the following sections, copied from a set of object or DWARF object files, and combined, section by section:

- .debug\_info.dwo
- .debug\_abbrev.dwo
- .debug\_line.dwo
- .debug\_loclists.dwo
- .debug\_rnglists.dwo
- .debug\_str\_offsets.dwo
- .debug\_str.dwo
- .debug\_macro.dwo

1 The string table section in `.debug_str.dwo` contains all the strings referenced  
2 from DWARF attributes using any of the forms `DW_FORM_strx`,  
3 `DW_FORM_strx1`, `DW_FORM_strx2`, `DW_FORM_strx3` or `DW_FORM_strx4`.  
4 Any attribute in a compilation unit or a type unit using this form refers to an  
5 entry in that unit's contribution to the `.debug_str_offsets.dwo` section, which  
6 in turn provides the offset of a string in the `.debug_str.dwo` section.

7 The DWARF package file also contains two index sections that provide a fast way  
8 to locate debug information by compilation unit ID for compilation units, or by  
9 type signature for type units:

```
10     .debug_cu_index  
11     .debug_tu_index
```

### 12 7.3.5.1 The Compilation Unit (CU) Index Section

13 The `.debug_cu_index` section is a hashed lookup table that maps a compilation  
14 unit ID to a set of contributions in the various debug information sections. Each  
15 contribution is stored as an offset within its corresponding section and a size.

16 Each compilation unit set may contain contributions from the following sections:

```
17     .debug_info.dwo (required)  
18     .debug_abbrev.dwo (required)  
19     .debug_line.dwo  
20     .debug_loclists.dwo  
21     .debug_rnglists.dwo  
22     .debug_str_offsets.dwo  
23     .debug_macro.dwo
```

24 *Note that a compilation unit set is not able to represent `.debug_macinfo` information*  
25 *from DWARF Version 4 or earlier formats.*

### 26 7.3.5.2 The Type Unit (TU) Index Section

27 The `.debug_tu_index` section is a hashed lookup table that maps a type signature  
28 to a set of offsets in the various debug information sections. Each contribution is  
29 stored as an offset within its corresponding section and a size.

30 Each type unit set may contain contributions from the following sections:

```
31     .debug_info.dwo (required)  
32     .debug_abbrev.dwo (required)  
33     .debug_line.dwo  
34     .debug_str_offsets.dwo
```

1 **7.3.5.3 Format of the CU and TU Index Sections**

2 Both `.debug_cu_index` and `.debug_tu_index` index sections have the same  
3 format, and serve to map an 8-byte signature to a set of contributions to the  
4 debug sections. Each index section begins with a header, followed by a hash table  
5 of signatures, a parallel table of indexes, a table of offsets, and a table of sizes.  
6 The index sections are aligned at 8-byte boundaries in the DWARF package file.

7 The index section header contains the following fields:

8 1. `version` (uhalf)

9 A version number. This number is specific to the CU and TU index  
10 information and is independent of the DWARF version number.

11 The version number is 6.

12 2. `offset_size_flag` (uhalf)

13 If the `offset_size_flag` is zero, the header is for a 32-bit DWARF format unit  
14 index section and all offsets and lengths are 4 bytes long; if it is one, the  
15 header is for a 64-bit DWARF format unit index section and all offsets and  
16 lengths are 8 bytes long.

17 3. `padding` (uhalf)

18 Reserved to DWARF (must be zero).

19 4. `section_count` (uword)

20 The number of entries in the table of section counts that follows. For brevity,  
21 the contents of this field is referred to as  $N$  below.

22 5. `unit_count` (uword)

23 The number of compilation units or type units in the index. For brevity, the  
24 contents of this field is referred to as  $U$  below.

25 6. `slot_count` (uword)

26 The number of slots in the hash table. For brevity, the contents of this field is  
27 referred to as  $S$  below.

28 *We assume that  $U$  and  $S$  do not exceed  $2^{32}$ .*

29 The size of the hash table,  $S$ , must be  $2^k$  such that:  $2^k > 3 * U/2$

30 The hash table begins at offset 16 in the section, and consists of an array of  $S$   
31 8-byte slots. Each slot contains a 64-bit signature.

32 The parallel table of indices begins immediately after the hash table (at offset  
33  $16 + 8 * S$  from the beginning of the section), and consists of an array of  $S$  4-byte  
34 slots, corresponding 1-1 with slots in the hash table. Each entry in the parallel  
35 table contains a row index into the tables of offsets and sizes.

## Chapter 7. Data Representation

1 Unused slots in the hash table have 0 in both the hash table entry and the parallel  
2 table entry. While 0 is a valid hash value, the row index in a used slot will always  
3 be non-zero.

4 Given an 8-byte compilation unit ID or type signature  $X$ , an entry in the hash  
5 table is located as follows:

- 6 1. Define  $REP(X)$  to be the value of  $X$  interpreted as an unsigned 64-bit integer  
7 in the target byte order.
- 8 2. Calculate a primary hash  $H = REP(X) \& MASK(k)$ , where  $MASK(k)$  is a  
9 mask with the low-order  $k$  bits all set to 1.
- 10 3. Calculate a secondary hash  $H' = (((REP(X) \gg 32) \& MASK(k)) | 1)$ .
- 11 4. If the hash table entry at index  $H$  matches the signature, use that entry. If the  
12 hash table entry at index  $H$  is unused (all zeroes), terminate the search: the  
13 signature is not present in the table.
- 14 5. Let  $H = (H + H') \text{ modulo } S$ . Repeat at Step 4.

15 Because  $S > U$ , and  $H'$  and  $S$  are relatively prime, the search is guaranteed to  
16 stop at an unused slot or find the match.

17 The table of offsets begins immediately following the parallel table (at offset  
18  $16 + 12 * S$  from the beginning of the section). This table consists of a single  
19 header row containing  $N$  fields, each a 4-byte unsigned integer, followed by  $U$   
20 data rows, each containing  $N$  unsigned integer fields of size specified by the  
21 index header `offset_size_flag` field. The fields in the header row provide a  
22 section identifier referring to a debug section; the available section identifiers are  
23 shown in Table 7.1 following. Each data row corresponds to a specific CU or TU  
24 in the package file. In the data rows, each field provides an offset to the debug  
25 section whose identifier appears in the corresponding field of the header row.  
26 The data rows are indexed starting at 1.

27 *Not all sections listed in the table need be included.*

## Chapter 7. Data Representation

Table 7.1: DWARF package file section identifier encodings

Section identifier	Value	Section
DW_SECT_INFO	1	.debug_info.dwo
<i>Reserved</i>	2	
DW_SECT_ABBREV	3	.debug_abbrev.dwo
DW_SECT_LINE	4	.debug_line.dwo
DW_SECT_LOCLISTS	5	.debug_loclists.dwo
DW_SECT_STR_OFFSETS	6	.debug_str_offsets.dwo
DW_SECT_MACRO	7	.debug_macro.dwo
DW_SECT_RNGLISTS	8	.debug_rnglists.dwo

1 The offsets provided by the CU and TU index sections are the base offsets for the  
2 contributions made by each CU or TU to the corresponding section in the  
3 package file. Each CU and TU header contains a `debug_abbrev_offset` field,  
4 used to find the abbreviations table for that CU or TU within the contribution to  
5 the `.debug_abbrev.dwo` section for that CU or TU, and are interpreted as relative  
6 to the base offset given in the index section. Likewise, offsets into  
7 `.debug_line.dwo` from [DW\\_AT\\_stmt\\_list](#) attributes are interpreted as relative to  
8 the base offset for `.debug_line.dwo`, and offsets into other debug sections  
9 obtained from DWARF attributes are also interpreted as relative to the  
10 corresponding base offset.

11 The table of sizes begins immediately following the table of offsets, and provides  
12 the sizes of the contributions made by each CU or TU to the corresponding  
13 section in the package file. This table consists of  $U$  data rows, each with  $N$   
14 unsigned integer fields of size specified by the index header `offset_size_flag`  
15 field. Each data row corresponds to the same CU or TU as the corresponding  
16 data row in the table of offsets described above. Within each data row, the  $N$   
17 fields also correspond one-to-one with the fields in the corresponding data row  
18 of the table of offsets. Each field provides the size of the contribution made by a  
19 CU or TU to the corresponding section in the package file.

20 For an example, see [Figure F.10 on page 433](#).



### 7.3.6 DWARF Supplementary Object Files

*A supplementary object file permits a post-link utility to analyze executable and shared object files and collect duplicate debugging information into a single file that can be referenced by each of the original files. This is in contrast to split DWARF object files, which allow the compiler to split the debugging information between multiple files in order to reduce link time and executable size.*

A DWARF supplementary object file is itself an object file, using the same object file format, byte order, and size as the corresponding application executables or shared libraries. It contains a file header, section table, and a number of DWARF debug information sections. Both the supplementary object file and all the executable or shared object files that reference entries or strings in that file must contain a `.debug_sup` section that establishes the relationship.

The `.debug_sup` section contains:

1. `version` (uhalf)

A 2-byte unsigned integer representing the version of the DWARF information for the compilation unit.

The value in this field is 5.

2. `is_supplementary` (ubyte)

A 1-byte unsigned integer, which contains the value 1 if it is in the supplementary object file that other executable or shared object files refer to, or 0 if it is an executable or shared object referring to a supplementary object file.

3. `sup_filename` (null terminated filename string)

If `is_supplementary` is 0, this contains either an absolute filename for the supplementary object file, or a filename relative to the object file containing the `.debug_sup` section. If `is_supplementary` is 1, then `sup_filename` is not needed and must be an empty string (a single null byte).

4. `sup_checksum_len` (unsigned LEB128)

Length of the following `sup_checksum` field; this value can be 0 if no checksum is provided.

5. `sup_checksum` (array of ubyte)

An implementation-defined integer constant value that provides unique identification of the supplementary file.

1 Debug information entries that refer to an executable's or shared object's  
2 addresses must *not* be moved to supplementary files (the addresses will likely not  
3 be the same). Similarly, entries referenced from within location descriptions or  
4 using `loclistsptr` form attributes must not be moved to a supplementary object  
5 file.

6 Executable or shared object file compilation units can use  
7 `DW_TAG_imported_unit` with an `DW_AT_import` attribute that uses form  
8 `DW_FORM_ref_sup4` or `DW_FORM_ref_sup8` to import entries from the  
9 supplementary object file, form `DW_FORM_ref_sup4` or `DW_FORM_ref_sup8` to  
10 refer directly to individual entries in the supplementary file, or form  
11 `DW_FORM_strp_sup` or `DW_FORM_strp_sup8` to refer to strings that are used  
12 by debug information of multiple executables or shared object files. Within the  
13 supplementary object file's debugging sections, forms `DW_FORM_ref_sup4`,  
14 `DW_FORM_ref_sup8`, `DW_FORM_strp_sup` and `DW_FORM_strp_sup8` are not  
15 used, and all reference forms referring to other sections refer to the local sections  
16 in the supplementary object file.

17 In macro information, `DW_MACRO_define_sup4`, `DW_MACRO_define_sup8`,  
18 `DW_MACRO_undef_sup4` and `DW_MACRO_undef_sup8` opcodes can refer to  
19 strings in the `.debug_str` section of the supplementary object file, while  
20 `DW_MACRO_import_sup4` and `DW_MACRO_import_sup8` can refer to  
21 `.debug_macro` section entries. Within the `.debug_macro` section of a  
22 supplementary object file, `DW_MACRO_define_strp` and  
23 `DW_MACRO_undef_strp` opcodes refer to the local `.debug_str` section in that  
24 supplementary file, not the one in the executable or shared object file.

25 *Forms for both 4- and 8-byte references are provided so that references may use the*  
26 *appropriate offset size for the content of the supplementary object file, which might not*  
27 *use the same 32-bit or 64-bit DWARF format as a referencing object file.*

## 28 7.4 32-Bit and 64-Bit DWARF Formats

29 There are two closely-related DWARF formats. In the 32-bit DWARF format, all  
30 values that represent lengths of DWARF sections and offsets relative to the  
31 beginning of DWARF sections are represented using four bytes. In the 64-bit  
32 DWARF format, all values that represent lengths of DWARF sections and offsets  
33 relative to the beginning of DWARF sections are represented using eight bytes. A  
34 special convention applies to the initial length field of certain DWARF sections,  
35 as well as the CIE and FDE structures, so that the 32-bit and 64-bit DWARF  
36 formats can coexist and be distinguished within a single linked object.

## Chapter 7. Data Representation

1 The 32-bit and 64-bit DWARF format conventions must *not* be intermixed within  
2 a single compilation unit, except for contributions to the `.debug_str_offsets` (or  
3 `.debug_str_offsets.dwo`) section.

4 *The exception for the `.debug_str_offsets` section enables an executable program with*  
5 *a mixture of 32-bit and 64-bit DWARF compilation units to refer to any string in the*  
6 *merged `.debug_str` section, even if that section exceeds 4GB in size.*

7 Except where noted otherwise, all references in this document to a debugging  
8 information section (for example, `.debug_info`), apply also to the corresponding  
9 split DWARF section (for example, `.debug_info.dwo`).

10 *Attribute values and section header fields that represent addresses in the target program*  
11 *are not affected by the rules that follow.*

## Chapter 7. Data Representation

1 The differences between the 32- and 64-bit DWARF formats are detailed in the  
2 following:

- 3 1. In the 32-bit DWARF format, an initial length field (see Section 7.2.2 on  
4 [page 192](#)). is an unsigned 4-byte integer (which must be less than  
5 0xffffffff0); in the 64-bit DWARF format, an initial length field is 12 bytes in  
6 size, and has two parts:

- 7 • The first four bytes have the value 0xffffffff.  
8 • The following eight bytes contain the actual length represented as an  
9 unsigned 8-byte integer.

10 *This representation allows a DWARF consumer to dynamically detect that a*  
11 *DWARF section contribution is using the 64-bit format and to adapt its processing*  
12 *accordingly.*

- 13 2. Section offset and section length fields that occur in the headers of DWARF  
14 sections (other than initial length fields) depend on the choice of DWARF  
15 format as follows: for the 32-bit DWARF format these are 4-byte unsigned  
16 integer values; for the 64-bit DWARF format, they are 8-byte unsigned integer  
17 values.

Section	Name	Role
.debug_aranges	debug_info_offset	offset in .debug_info
.debug_frame/CIE	CIE_id	CIE distinguished value
.debug_frame/FDE	CIE_pointer	offset in .debug_frame
.debug_info	debug_abbrev_offset	offset in .debug_abbrev
.debug_line	header_length	length of header itself
.debug_names	entry in array of CUs or local TUs	offset in .debug_info

## Chapter 7. Data Representation

1 The CIE\_id field in a CIE structure must be 64 bits because it overlays the  
2 CIE\_pointer in a FDE structure; this implicit union must be accessed to  
3 distinguish whether a CIE or FDE is present, consequently, these two fields  
4 must exactly overlay each other (both offset and size).

- 5 3. Within the body of the .debug\_info section, certain forms of attribute value  
6 depend on the choice of DWARF format as follows: for the 32-bit DWARF  
7 format, the value is a 4-byte unsigned integer; for the 64-bit DWARF format,  
8 the value is an 8-byte unsigned integer.

Form	Role
DW_FORM_line_strp	offset in .debug_line_str
DW_FORM_ref_addr	offset in .debug_info
DW_FORM_sec_offset	offset in a section other than .debug_info or .debug_str
DW_FORM_strp	offset in .debug_str
DW_FORM_strp_sup	offset in .debug_str section of a supplementary object file
DW_OP_call_ref	offset in .debug_info

- 9 4. Within the body of the .debug\_line section, certain forms of content  
10 description depend on the choice of DWARF format as follows: for the 32-bit  
11 DWARF format, the value is a 4-byte unsigned integer; for the 64-bit DWARF  
12 format, the value is a 8-byte unsigned integer.

Form	Role
DW_FORM_line_strp	offset in .debug_line_str

- 13 5. Within the body of the .debug\_names sections, the representation of each  
14 entry in the array of compilation units (CUs) and the array of local type units  
15 (TUs), which represents an offset in the .debug\_info section, depends on the  
16 DWARF format as follows: for the 32-bit DWARF format, each entry is a  
17 4-byte unsigned integer; for the 64-bit DWARF format, it is a 8-byte unsigned  
18 integer.
- 19 6. In the body of the .debug\_str\_offsets sections, the size of entries in the  
20 body depend on the DWARF format as follows: for the 32-bit DWARF format,  
21 entries are 4-byte unsigned integer values; for the 64-bit DWARF format, they  
22 are 8-byte unsigned integers.

1 7. Within the body of the `.debug_loclists` and `.debug_rnglists` sections, the  
2 offsets that follow the header depend on the DWARF format as follows: for  
3 the 32-bit DWARF format, offsets are 4-byte unsigned integer values; for the  
4 64-bit DWARF format, they are 8-byte unsigned integers.

5 A DWARF consumer that supports the 64-bit DWARF format must support  
6 executables in which some compilation units use the 32-bit format and others use  
7 the 64-bit format provided that the combination links correctly (that is, provided  
8 that there are no link-time errors due to truncation or overflow). (An  
9 implementation is not required to guarantee detection and reporting of all such  
10 errors.)

11 *It is expected that DWARF producing compilers will not use the 64-bit format by*  
12 *default. In most cases, the division of even very large applications into a number of*  
13 *executable and shared object files will suffice to assure that the DWARF sections within*  
14 *each individual linked object are less than 4 GBytes in size. However, for those cases*  
15 *where needed, the 64-bit format allows the unusual case to be handled as well. Even in*  
16 *this case, it is expected that only application supplied objects will need to be compiled*  
17 *using the 64-bit format; separate 32-bit format versions of system supplied shared*  
18 *executable libraries can still be used.*

## 19 7.5 Format of Debugging Information

20 For each compilation unit compiled with a DWARF producer, a contribution is  
21 made to the `.debug_info` section of the object file. Each such contribution  
22 consists of a compilation unit header (see Section [7.5.1.1 on the following page](#))  
23 followed by a single `DW_TAG_compile_unit` or `DW_TAG_partial_unit`  
24 debugging information entry, together with its children.

25 For each type defined in a compilation unit, a separate contribution may also be  
26 made to the `.debug_info` section of the object file. Each such contribution  
27 consists of a type unit header (see Section [7.5.1.3 on page 211](#)) followed by a  
28 `DW_TAG_type_unit` entry, together with its children.

29 Each debugging information entry begins with a code that represents an entry in  
30 a separate abbreviations table. This code is followed directly by a series of  
31 attribute values.

32 The appropriate entry in the abbreviations table guides the interpretation of the  
33 information contained directly in the `.debug_info` section.

34 Multiple debugging information entries may share the same abbreviation table  
35 entry. Each compilation unit is associated with a particular abbreviation table,  
36 but multiple compilation units may share the same table.

## 7.5.1 Unit Headers

Unit headers contain a field, `unit_type`, whose value indicates the kind of compilation unit (see Section 3.1 on page 61). The encodings for the unit type enumeration are shown in Table 7.2.

Table 7.2: Unit header unit type encodings

Unit header unit type encodings	Value
<code>DW_UT_compile</code>	0x01
<code>DW_UT_type</code>	0x02
<code>DW_UT_partial</code>	0x03
<code>DW_UT_skeleton</code>	0x04
<code>DW_UT_split_compile</code>	0x05
<code>DW_UT_split_type</code>	0x06
<code>DW_UT_lo_user</code>	0x80
<code>DW_UT_hi_user</code>	0xff

All unit headers have the same initial three fields: `initial_length`, `version` and `unit_type`.

### 7.5.1.1 Full and Partial Compilation Unit Headers

1. `unit_length` ([initial length](#))  
A 4-byte or 12-byte unsigned integer representing the length of the `.debug_info` contribution for that compilation unit, not including the length field itself (see Section 7.4 on page 204).

2. `version` (uhalf)  
A 2-byte unsigned integer representing the version of the DWARF information for the compilation unit.

The value in this field is 5.

*See also Appendix G on page 436 for a summary of all version numbers that apply to DWARF sections.*

1 3. `unit_type` (ubyte)

2 A 1-byte unsigned integer identifying this unit as a compilation unit. The  
3 value of this field is [DW\\_UT\\_compile](#) for a (non-split) full compilation unit or  
4 [DW\\_UT\\_partial](#) for a (non-split) partial compilation unit (see Section 3.1.1 on  
5 page 62).

6 *See Section 7.5.1.2 regarding a split full compilation unit.*

7 4. `address_size` (ubyte)

8 A 1-byte unsigned integer representing the size in bytes of an address on the  
9 target architecture.

10 5. `debug_abbrev_offset` (section offset)

11 A 4-byte or 8-byte unsigned offset into the `.debug_abbrev` section. This offset  
12 associates the compilation unit with a particular set of debugging  
13 information entry abbreviations. In the [32-bit DWARF format](#), this is a 4-byte  
14 unsigned length; in the [64-bit DWARF format](#), this is an 8-byte unsigned  
15 length (see Section 7.4 on page 204).

16 **7.5.1.2 Skeleton and Split Compilation Unit Headers**

17 1. `unit_length` (initial length)

18 A 4-byte or 12-byte unsigned integer representing the length of the  
19 `.debug_info` contribution for that compilation unit, not including the length  
20 field itself (see Section 7.4 on page 204).

21 2. `version` (uhalf)

22 A 2-byte unsigned integer representing the version of the DWARF  
23 information for the compilation unit.

24 The value in this field is 5.

25 *See also Appendix G on page 436 for a summary of all version numbers that apply to*  
26 *DWARF sections.*

27 3. `unit_type` (ubyte)

28 A 1-byte unsigned integer identifying this unit as a compilation unit. The  
29 value of this field is [DW\\_UT\\_skeleton](#) for a skeleton compilation unit or  
30 [DW\\_UT\\_split\\_compile](#) for a split (full) compilation unit (see Section 3.1.2 on  
31 page 69).

32 *There is no split analog to the partial compilation unit.*



- 1 4. `address_size` (ubyte)  
2 A 1-byte unsigned integer representing the size in bytes of an address on the  
3 target architecture.
- 4 5. `debug_abbrev_offset` ([section offset](#))  
5 A 4-byte or 8-byte unsigned offset into the `.debug_abbrev` section. This offset  
6 associates the compilation unit with a particular set of debugging  
7 information entry abbreviations. In the [32-bit DWARF format](#), this is a 4-byte  
8 unsigned length; in the [64-bit DWARF format](#), this is an 8-byte unsigned  
9 length (see [Section 7.4 on page 204](#)).
- 10 6. `dwo_id` (unit ID)  
11 An 8-byte implementation-defined integer constant value, known as the  
12 compilation unit ID, that provides unique identification of a skeleton  
13 compilation unit and its associated split compilation unit in the object file  
14 named in the [DW\\_AT\\_dwo\\_name](#) attribute of the skeleton compilation.

### 15 7.5.1.3 Type Unit Headers

16 The header for the series of debugging information entries contributing to the  
17 description of a type that has been placed in its own type unit, within the  
18 `.debug_info` section, consists of the following information:

- 19 1. `unit_length` ([initial length](#))  
20 A 4-byte or 12-byte unsigned integer representing the length of the  
21 `.debug_info` contribution for that type unit, not including the length field  
22 itself (see [Section 7.4 on page 204](#)).
- 23 2. `version` (uhalf)  
24 A 2-byte unsigned integer representing the version of the DWARF  
25 information for the type unit.  
26 The value in this field is 5.
- 27 3. `unit_type` (ubyte)  
28 A 1-byte unsigned integer identifying this unit as a type unit. The value of  
29 this field is [DW\\_UT\\_type](#) for a non-split type unit (see [Section 3.1.4 on](#)  
30 [page 71](#)) or [DW\\_UT\\_split\\_type](#) for a split type unit.
- 31 4. `address_size` (ubyte)  
32 A 1-byte unsigned integer representing the size in bytes of an address on the  
33 target architecture.

1 5. `debug_abbrev_offset` (section offset)

2 A 4-byte or 8-byte unsigned offset into the `.debug_abbrev` section. This offset  
3 associates the type unit with a particular set of debugging information entry  
4 abbreviations. In the [32-bit DWARF format](#), this is a 4-byte unsigned length;  
5 in the [64-bit DWARF format](#), this is an 8-byte unsigned length (see Section 7.4  
6 on page 204).

7 6. `type_signature` (8-byte unsigned integer)

8 A unique 8-byte signature (see Section 7.32 on page 257) of the type described  
9 in this type unit.

10 *An attribute that refers (using `DW_FORM_ref_sig8`) to the primary type contained*  
11 *in this type unit uses this value.*

12 7. `type_offset` (section offset)

13 A 4-byte or 8-byte unsigned offset relative to the beginning of the type unit  
14 header. This offset refers to the debugging information entry that describes  
15 the type. Because the type may be nested inside a namespace or other  
16 structures, and may contain references to other types that have not been  
17 placed in separate type units, it is not necessarily either the first or the only  
18 entry in the type unit. In the [32-bit DWARF format](#), this is a 4-byte unsigned  
19 length; in the [64-bit DWARF format](#), this is an 8-byte unsigned length (see  
20 Section 7.4 on page 204).

21 **7.5.2 Debugging Information Entry**

22 Each debugging information entry begins with an unsigned LEB128 number  
23 containing the abbreviation code for the entry. This code represents an entry  
24 within the abbreviations table associated with the compilation unit containing  
25 this entry. The abbreviation code is followed by a series of attribute values.

26 On some architectures, there are alignment constraints on section boundaries. To  
27 make it easier to pad debugging information sections to satisfy such constraints,  
28 the abbreviation code 0 is reserved. Debugging information entries consisting of  
29 only the abbreviation code 0 are considered null entries.

1 **7.5.3 Abbreviations Tables**

2 The abbreviations tables for all compilation units are contained in a separate  
3 object file section called `.debug_abbrev`. As mentioned before, multiple  
4 compilation units may share the same abbreviations table.

5 The abbreviations table for a single compilation unit consists of a series of  
6 abbreviation declarations. Each declaration specifies the tag and attributes for a  
7 particular form of debugging information entry. Each declaration begins with an  
8 unsigned LEB128 number representing the abbreviation code itself. It is this code  
9 that appears at the beginning of a debugging information entry in the  
10 `.debug_info` section. As described above, the abbreviation code 0 is reserved for  
11 null debugging information entries. The abbreviation code is followed by  
12 another unsigned LEB128 number that encodes the entry's tag. The encodings  
13 for the tag names are given in Table 7.3 on the next page.

*Continued on next page*

## Chapter 7. Data Representation

Tag name	Value
----------	-------

Table 7.3: Tag encodings

Tag name	Value
DW_TAG_array_type	0x01
DW_TAG_class_type	0x02
DW_TAG_entry_point	0x03
DW_TAG_enumeration_type	0x04
DW_TAG_formal_parameter	0x05
<i>Reserved</i>	0x06
<i>Reserved</i>	0x07
DW_TAG_imported_declaration	0x08
<i>Reserved</i>	0x09
DW_TAG_label	0x0a
DW_TAG_lexical_block	0x0b
<i>Reserved</i>	0x0c
DW_TAG_member	0x0d
<i>Reserved</i>	0x0e
DW_TAG_pointer_type	0x0f
DW_TAG_reference_type	0x10
DW_TAG_compile_unit	0x11
DW_TAG_string_type	0x12
DW_TAG_structure_type	0x13
<i>Reserved</i>	0x14
DW_TAG_subroutine_type	0x15
DW_TAG_typedef	0x16
DW_TAG_union_type	0x17
DW_TAG_unspecified_parameters	0x18
DW_TAG_variant	0x19
DW_TAG_common_block	0x1a
DW_TAG_common_inclusion	0x1b
DW_TAG_inheritance	0x1c
DW_TAG_inlined_subroutine	0x1d

*Continued on next page*

## Chapter 7. Data Representation

Tag name	Value
DW_TAG_module	0x1e
DW_TAG_ptr_to_member_type	0x1f
DW_TAG_set_type	0x20
DW_TAG_subrange_type	0x21
DW_TAG_with_stmt	0x22
DW_TAG_access_declaration	0x23
DW_TAG_base_type	0x24
DW_TAG_catch_block	0x25
DW_TAG_const_type	0x26
DW_TAG_constant	0x27
DW_TAG_enumerator	0x28
DW_TAG_file_type	0x29
DW_TAG_friend	0x2a
DW_TAG_namelist	0x2b
DW_TAG_namelist_item	0x2c
DW_TAG_packed_type	0x2d
DW_TAG_subprogram	0x2e
DW_TAG_template_type_parameter	0x2f
DW_TAG_template_value_parameter	0x30
DW_TAG_thrown_type	0x31
DW_TAG_try_block	0x32
DW_TAG_variant_part	0x33
DW_TAG_variable	0x34
DW_TAG_volatile_type	0x35
DW_TAG_dwarf_procedure	0x36
DW_TAG_restrict_type	0x37
DW_TAG_interface_type	0x38
DW_TAG_namespace	0x39
DW_TAG_imported_module	0x3a
DW_TAG_unspecified_type	0x3b
DW_TAG_partial_unit	0x3c
DW_TAG_imported_unit	0x3d

*Continued on next page*

## Chapter 7. Data Representation

Tag name	Value
<i>Reserved</i>	0x3e <sup>1</sup>
DW_TAG_condition	0x3f
DW_TAG_shared_type	0x40
DW_TAG_type_unit	0x41
DW_TAG_rvalue_reference_type	0x42
DW_TAG_template_alias	0x43
DW_TAG_coarray_type	0x44
DW_TAG_generic_subrange	0x45
DW_TAG_dynamic_type	0x46
DW_TAG_atomic_type	0x47
DW_TAG_call_site	0x48
DW_TAG_call_site_parameter	0x49
DW_TAG_skeleton_unit	0x4a
DW_TAG_immutable_type	0x4b
DW_TAG_lo_user	0x4080
DW_TAG_hi_user	0xffff

1 Following the tag encoding is a 1-byte value that determines whether a  
2 debugging information entry using this abbreviation has child entries or not. If  
3 the value is **DW\_CHILDREN\_yes**, the next physically succeeding entry of any  
4 debugging information entry using this abbreviation is the first child of that  
5 entry. If the 1-byte value following the abbreviation's tag encoding is  
6 **DW\_CHILDREN\_no**, the next physically succeeding entry of any debugging  
7 information entry using this abbreviation is a sibling of that entry. (Either the  
8 first child or sibling entries may be null entries). The encodings for the child  
9 determination byte are given in Table 7.4 on the following page (As mentioned in  
10 Section 2.3 on page 25, each chain of sibling entries is terminated by a null entry.)

---

<sup>1</sup>Code 0x3e is reserved to allow backward compatible support of the DW\_TAG\_mutable\_type DIE that was defined (only) in DWARF Version 3.

## Chapter 7. Data Representation

Table 7.4: Child determination encodings

Children determination name	Value
DW_CHILDREN_no	0x00
DW_CHILDREN_yes	0x01

1 Finally, the child encoding is followed by a series of attribute specifications. Each  
2 attribute specification consists of two parts (except for  
3 [DW\\_FORM\\_implicit\\_const](#), [DW\\_FORM\\_addrx\\_offset](#) and [DW\\_FORM\\_indirect](#),  
4 see below). The first part is an unsigned LEB128 number representing the  
5 attribute's name. The second part is an unsigned LEB128 number representing  
6 the attribute's form. The series of attribute specifications ends with an entry  
7 containing 0 for the name and 0 for the form.

8 For attributes with the form [DW\\_FORM\\_implicit\\_const](#), in addition to the  
9 attribute name and form values, the attribute specification contains a third part,  
10 which is a signed LEB128 number. The value of this number is used as the value  
11 of the attribute.

12 For attributes with the form [DW\\_FORM\\_addrx\\_offset](#), following the attribute  
13 name, the attribute specification contains two unsigned LEB128 numbers, each  
14 representing a form. The first form must be of class address and the second of  
15 class constant. Values using this form in the `.debug_info` section contain a value  
16 for the first form followed by a value for the second form. The total value of the  
17 [DW\\_FORM\\_addrx\\_offset](#) is then computed by adding those two values together  
18 (if the first value is an indirect address, that is resolved first before adding it to  
19 the second value).

20 For attributes with the form [DW\\_FORM\\_indirect](#), the actual attribute form value  
21 itself is in the `.debug_info` section which begins with an unsigned LEB128  
22 number that specifies the actual form, followed by the value according to that  
23 form. This allows producers to choose forms for particular attributes  
24 dynamically, without having to add a new entry to the abbreviations table.

25 If the actual attribute form is [DW\\_FORM\\_implicit\\_const](#), the form is (still)  
26 followed by a signed LEB128 number.

27 If the actual attribute form is itself [DW\\_FORM\\_indirect](#), the indirection repeats.  
28 There may be one or more occurrences of [DW\\_FORM\\_indirect](#) in sequence until  
29 a non-[DW\\_FORM\\_indirect](#) form is reached. The sequence of  
30 [DW\\_FORM\\_indirect](#) forms does not have any effect other than to use up space.

31 The abbreviations for a given compilation unit end with an entry consisting of a  
32 0 byte for the abbreviation code.

1 See Appendix D.1.1 on page 298 for a depiction of the organization of the debugging  
2 information.

### 3 7.5.4 Attribute Encodings

4 The encodings for the attribute names are given in Table 7.5 following.

Table 7.5: Attribute encodings

Attribute name	Value	Classes
DW_AT_sibling	0x01	reference
DW_AT_location	0x02	locdesc, loclist
DW_AT_name	0x03	string
Reserved	0x04	not applicable
Reserved	0x05	not applicable
Reserved	0x06	not applicable
Reserved	0x07	not applicable
Reserved	0x08	not applicable
DW_AT_ordering	0x09	constant
Reserved	0x0a	not applicable
DW_AT_byte_size	0x0b	constant, exprval, reference
Reserved	0x0c <sup>2</sup>	not applicable
DW_AT_bit_size	0x0d	constant, exprval, reference
Reserved	0x0e	not applicable
Reserved	0x0f	not applicable
DW_AT_stmt_list	0x10	lineptr
DW_AT_low_pc	0x11	address
DW_AT_high_pc	0x12	address, constant
Reserved	0x13 <sup>3</sup>	not applicable
Reserved	0x14	not applicable
DW_AT_discr	0x15	reference
DW_AT_discr_value	0x16	constant

*Continued on next page*

<sup>2</sup>Code 0x0c is reserved to allow backward compatible support of the DW\_AT\_bit\_offset attribute which was defined in DWARF Version 3 and earlier.

<sup>3</sup>Code 0x13 is reserved to allow backward compatible support of the DW\_AT\_language attribute which was defined in DWARF Version 5 and earlier.



## Chapter 7. Data Representation

Attribute name	Value	Classes
DW_AT_visibility	0x17	constant
DW_AT_import	0x18	reference
DW_AT_string_length	0x19	locdesc, loclist, reference
DW_AT_common_reference	0x1a	reference
DW_AT_comp_dir	0x1b	string
DW_AT_const_value	0x1c	block, constant, string
DW_AT_containing_type	0x1d	reference
DW_AT_default_value	0x1e	constant, reference, flag, string
<i>Reserved</i>	0x1f	<i>not applicable</i>
DW_AT_inline	0x20	constant
DW_AT_is_optional	0x21	flag
DW_AT_lower_bound	0x22	constant, exprval, reference
<i>Reserved</i>	0x23	<i>not applicable</i>
<i>Reserved</i>	0x24	<i>not applicable</i>
DW_AT_producer	0x25	string
<i>Reserved</i>	0x26	<i>not applicable</i>
DW_AT_prototyped	0x27	flag
<i>Reserved</i>	0x28	<i>not applicable</i>
<i>Reserved</i>	0x29	<i>not applicable</i>
DW_AT_return_addr	0x2a	locdesc, loclist
<i>Reserved</i>	0x2b	<i>not applicable</i>
DW_AT_start_scope	0x2c	constant, rnglist
<i>Reserved</i>	0x2d	<i>not applicable</i>
DW_AT_bit_stride	0x2e	constant, exprval, reference
DW_AT_upper_bound	0x2f	constant, exprval, reference
<i>Reserved</i>	0x30	<i>not applicable</i>
DW_AT_abstract_origin	0x31	reference
DW_AT_accessibility	0x32	constant
DW_AT_address_class	0x33	constant
DW_AT_artificial	0x34	flag
DW_AT_base_types	0x35	reference
DW_AT_calling_convention	0x36	constant

*Continued on next page*

## Chapter 7. Data Representation

Attribute name	Value	Classes
DW_AT_count	0x37	constant, exprval, reference
DW_AT_data_member_location	0x38	constant, locdesc, loclist
DW_AT_decl_column	0x39	constant
DW_AT_decl_file	0x3a	constant
DW_AT_decl_line	0x3b	constant
DW_AT_declaration	0x3c	flag
DW_AT_discr_list	0x3d	block
DW_AT_encoding	0x3e	constant
DW_AT_external	0x3f	flag
DW_AT_frame_base	0x40	locdesc, loclist
DW_AT_friend	0x41	reference
DW_AT_identifier_case	0x42	constant
<i>Reserved</i>	0x43 <sup>4</sup>	<i>not applicable</i>
DW_AT_namelist_item	0x44	reference
DW_AT_priority	0x45	reference
<i>Reserved</i>	0x46 <sup>5</sup>	<i>not applicable</i>
DW_AT_specification	0x47	reference
DW_AT_static_link	0x48	locdesc, loclist
DW_AT_type	0x49	reference
DW_AT_use_location	0x4a	locdesc, loclist
DW_AT_variable_parameter	0x4b	flag
DW_AT_virtuality	0x4c	constant
DW_AT_vtable_elem_location	0x4d	locdesc, loclist
DW_AT_allocated	0x4e	constant, exprval, reference
DW_AT_associated	0x4f	constant, exprval, reference
DW_AT_data_location	0x50	locdesc
DW_AT_byte_stride	0x51	constant, exprval, reference
DW_AT_entry_pc	0x52	address, constant
DW_AT_use_UTF8	0x53	flag

*Continued on next page*

<sup>4</sup>Code 0x43 is reserved to allow backward compatible support of the DW\_AT\_macro\_info attribute which was defined in DWARF Version 4 and earlier.

<sup>5</sup>Code 0x46 is reserved to allow backward compatible support of the DW\_AT\_segment attribute which was defined in DWARF Version 5 and earlier.

## Chapter 7. Data Representation

Attribute name	Value	Classes
DW_AT_extension	0x54	reference
DW_AT_ranges	0x55	rnglist
DW_AT_trampoline	0x56	address, flag, reference, string
DW_AT_call_column	0x57	constant
DW_AT_call_file	0x58	constant
DW_AT_call_line	0x59	constant
DW_AT_description	0x5a	string
DW_AT_binary_scale	0x5b	constant
DW_AT_decimal_scale	0x5c	constant
DW_AT_small	0x5d	reference
DW_AT_decimal_sign	0x5e	constant
DW_AT_digit_count	0x5f	constant
DW_AT_picture_string	0x60	string
DW_AT_mutable	0x61	flag
DW_AT_threads_scaled	0x62	flag
DW_AT_explicit	0x63	flag
DW_AT_object_pointer	0x64	reference
DW_AT_endianity	0x65	constant
DW_AT_elemental	0x66	flag
DW_AT_pure	0x67	flag
DW_AT_recursive	0x68	flag
DW_AT_signature	0x69	reference
DW_AT_main_subprogram	0x6a	flag
DW_AT_data_bit_offset	0x6b	constant
DW_AT_const_expr	0x6c	flag
DW_AT_enum_class	0x6d	flag
DW_AT_linkage_name	0x6e	string
DW_AT_string_length_bit_size	0x6f	constant
DW_AT_string_length_byte_size	0x70	constant
DW_AT_rank	0x71	constant, exprval

*Continued on next page*

## Chapter 7. Data Representation

Attribute name	Value	Classes
<i>Reserved</i>	0x72 <sup>6</sup>	<i>not applicable</i>
DW_AT_addr_base	0x73	addrptr
DW_AT_rnglists_base	0x74	rnglistsptr
<i>Reserved</i>	0x75	<i>not applicable</i>
DW_AT_dwo_name	0x76	string
DW_AT_reference	0x77	flag
DW_AT_rvalue_reference	0x78	flag
DW_AT_macros	0x79	macptr
DW_AT_call_all_calls	0x7a	flag
DW_AT_call_all_source_calls	0x7b	flag
DW_AT_call_all_tail_calls	0x7c	flag
DW_AT_call_return_pc	0x7d	address
DW_AT_call_value	0x7e	exprval
DW_AT_call_origin	0x7f	reference
DW_AT_call_parameter	0x80	reference
DW_AT_call_pc	0x81	address
DW_AT_call_tail_call	0x82	flag
DW_AT_call_target	0x83	locdesc
DW_AT_call_target_clobbered	0x84	locdesc
DW_AT_call_data_location	0x85	locdesc
DW_AT_call_data_value	0x86	exprval
DW_AT_noreturn	0x87	flag
DW_AT_alignment	0x88	constant
DW_AT_export_symbols	0x89	flag
DW_AT_deleted	0x8a	flag
DW_AT_defaulted	0x8b	constant
DW_AT_loclists_base	0x8c	loclistsptr
DW_AT_scale_multiplier ‡	0x8d	constant
DW_AT_scale_divisor ‡	0x8e	constant
DW_AT_str_offsets ‡	0x8f	stroffsetsptr

*Continued on next page*

---

<sup>6</sup>Code 0x72 is reserved to allow backward compatible support of the DW\_AT\_str\_offsets\_base attribute which was defined in DWARF Version 5 and earlier.

Attribute name	Value	Classes
<a href="#">DW_AT_language_name</a> ‡	0x90	constant
<a href="#">DW_AT_language_version</a> ‡	0x91	constant
<a href="#">DW_AT_bias</a> ‡	0x92	constant
<a href="#">DW_AT_tensor</a> ‡	0x93	flag
<a href="#">DW_AT_lo_user</a>	0x2000	—
<a href="#">DW_AT_hi_user</a>	0x3fff	—

‡ New in DWARF Version 6

## 7.5.5 Classes and Forms

Each class is a set of forms which have related representations and which are given a common interpretation according to the attribute in which the form is used. The attribute form governs how the value of an attribute is encoded. The classes and the forms they include are listed below.

Form [DW\\_FORM\\_sec\\_offset](#) is a member of more than one class, namely [addrptr](#), [lineptr](#), [loclist](#), [loclistsptr](#), [macptr](#), [rnglist](#), [rnglistsptr](#), and [stroffsetsptr](#); as a result, it is not possible for an attribute to allow more than one of these classes. The list of classes allowed by the applicable attribute in [Table 7.5 on page 218](#) determines the class of the form.

In the form descriptions that follow, some forms are said to depend in part on the value of an attribute of the [associated compilation unit](#):

- In the case of a split DWARF object file, the associated compilation unit is the skeleton compilation unit corresponding to the containing unit.
- Otherwise, the associated compilation unit is the containing unit.

Each possible form belongs to one or more of the following classes (see [Table 2.3 on page 23](#) for a summary of the purpose and general usage of each class):

- [address](#)  
Represented as either:
  - An object of appropriate size to hold an address on the target machine ([DW\\_FORM\\_addr](#)). The size is encoded in the compilation unit header (see [Section 7.5.1.1 on page 209](#)). This address is relocatable in a relocatable object file and is relocated in an executable file or shared object file.
  - An indirect index into a table of addresses (as described in the previous bullet) in the `.debug_addr` section ([DW\\_FORM\\_addrx](#),

## Chapter 7. Data Representation

`DW_FORM_addrx1`, `DW_FORM_addrx2`, `DW_FORM_addrx3` and `DW_FORM_addrx4`). The representation of a `DW_FORM_addrx` value is an unsigned LEB128 value, which is interpreted as a zero-based index into an array of addresses in the `.debug_addr` section. The representation of a `DW_FORM_addrx1`, `DW_FORM_addrx2`, `DW_FORM_addrx3` or `DW_FORM_addrx4` value is a 1-, 2-, 3- or 4-byte unsigned integer value, respectively, which is similarly interpreted. The index is relative to the value of the `DW_AT_addr_base` attribute of the associated compilation unit.

- `addrptr`

This is an offset into the `.debug_addr` section (`DW_FORM_sec_offset`). It consists of an offset from the beginning of the `.debug_addr` section to the beginning of the list of machine addresses information for the referencing entity. It is relocatable in a relocatable object file, and relocated in an executable or shared object file. In the 32-bit DWARF format, this offset is a 4-byte unsigned value; in the 64-bit DWARF format, it is an 8-byte unsigned value (see Section 7.4 on page 204).

- `block`

Blocks come in four forms:

- A 1-byte length followed by 0 to 255 contiguous information bytes (`DW_FORM_block1`).
- A 2-byte length followed by 0 to 65,535 contiguous information bytes (`DW_FORM_block2`).
- A 4-byte length followed by 0 to 4,294,967,295 contiguous information bytes (`DW_FORM_block4`).
- An unsigned LEB128 length followed by the number of bytes specified by the length (`DW_FORM_block`).

In all forms, the length is the number of information bytes that follow. The information bytes may contain any mixture of relocated (or relocatable) addresses, references to other debugging information entries or data bytes.

- `constant`

There are eight forms of constants. There are fixed length constant data forms for one-, two-, four-, eight- and sixteen-byte values (respectively, `DW_FORM_data1`, `DW_FORM_data2`, `DW_FORM_data4`, `DW_FORM_data8` and `DW_FORM_data16`). There are variable length constant data forms encoded using signed LEB128 numbers (`DW_FORM_sdata`) and unsigned LEB128 numbers (`DW_FORM_udata`).

## Chapter 7. Data Representation

1 There is also an implicit constant ([DW\\_FORM\\_implicit\\_const](#), see Section  
2 [7.5.3 on page 217](#)), whose value is provided as part of an abbreviation  
3 specification.

4 The data in [DW\\_FORM\\_data1](#), [DW\\_FORM\\_data2](#), [DW\\_FORM\\_data4](#),  
5 [DW\\_FORM\\_data8](#) and [DW\\_FORM\\_data16](#) can be anything. Depending on  
6 context, it may be a signed integer, an unsigned integer, a floating-point  
7 constant, or anything else. A consumer must use context to know how to  
8 interpret the bits, which if they are target machine data (such as an integer  
9 or floating-point constant) will be in target machine byte order.

10 *If one of the [DW\\_FORM\\_data<n>](#) forms is used to represent a signed or unsigned  
11 integer, it can be hard for a consumer to discover the context necessary to  
12 determine which interpretation is intended. Producers are therefore strongly  
13 encouraged to use [DW\\_FORM\\_sdata](#) or [DW\\_FORM\\_adata](#) for signed and  
14 unsigned integers respectively, rather than [DW\\_FORM\\_data<n>](#).*

- 15 • [exprval](#)

16 A DWARF expression that evaluates to a value (see Section [2.5 on page 26](#)).  
17 This is represented as an unsigned LEB128 length, followed by a byte  
18 sequence of the specified length ([DW\\_FORM\\_exprval](#)) containing the  
19 expression.

- 20 • [flag](#)

21 A flag is represented explicitly as a single byte of data ([DW\\_FORM\\_flag](#)) or  
22 implicitly ([DW\\_FORM\\_flag\\_present](#)). In the first case, if the flag has value  
23 zero, it indicates the absence of the attribute; if the flag has a non-zero  
24 value, it indicates the presence of the attribute. In the second case, the  
25 attribute is implicitly indicated as present, and no value is encoded in the  
26 debugging information entry itself.

- 27 • [lineptr](#)

28 This is an offset into the `.debug_line` or `.debug_line.dwo` section  
29 ([DW\\_FORM\\_sec\\_offset](#)). It consists of an offset from the beginning of the  
30 `.debug_line` section to the first byte of the data making up the line number  
31 list for the compilation unit. It is relocatable in a relocatable object file, and  
32 relocated in an executable or shared object file. In the [32-bit DWARF  
33 format](#), this offset is a 4-byte unsigned value; in the [64-bit DWARF format](#),  
34 it is an 8-byte unsigned value (see Section [7.4 on page 204](#)).

## Chapter 7. Data Representation

- **locdesc**

A DWARF location description (see Section 2.6 on page 38). This is represented as an unsigned LEB128 length, followed by a byte sequence of the specified length (**DW\_FORM\_locdesc**) containing the location description.

- **loclist**

This is represented as either:

- An index into the `.debug_loclists` section (**DW\_FORM\_loclistx**). The unsigned ULEB operand identifies an offset location relative to the base of that section (the location of the first offset in the section, not the first byte of the section). The contents of that location is then added to the base to determine the location of the target list of entries.
- An offset into the `.debug_loclists` section (**DW\_FORM\_sec\_offset**). The operand consists of a byte offset from the beginning of the `.debug_loclists` section. It is relocatable in a relocatable object file, and relocated in an executable or shared object file. In the **32-bit DWARF format**, this offset is a 4-byte unsigned value; in the **64-bit DWARF format**, it is an 8-byte unsigned value (see Section 7.4 on page 204).

- **loclistsptr**

This is an offset into the `.debug_loclists` section (**DW\_FORM\_sec\_offset**). The operand consists of a byte offset from the beginning of the `.debug_loclists` section. It is relocatable in a relocatable object file, and relocated in an executable or shared object file. In the **32-bit DWARF format**, this offset is a 4-byte unsigned value; in the **64-bit DWARF format**, it is an 8-byte unsigned value (see Section 7.4 on page 204).

- **macptr**

This is an offset into the `.debug_macro` or `.debug_macro.dwo` section (**DW\_FORM\_sec\_offset**). It consists of an offset from the beginning of the `.debug_macro` or `.debug_macro.dwo` section to the the header making up the macro information list for the compilation unit. It is relocatable in a relocatable object file, and relocated in an executable or shared object file. In the **32-bit DWARF format**, this offset is a 4-byte unsigned value; in the **64-bit DWARF format**, it is an 8-byte unsigned value (see Section 7.4 on page 204).



1 • [rnglist](#)

2 This is represented as either:

- 3 – An index into the `.debug_rnglists` section ([DW\\_FORM\\_rnglistx](#)). The  
4 unsigned ULEB operand identifies an offset location relative to the  
5 base of that section (the location of the first offset in the section, not the  
6 first byte of the section). The contents of that location is then added to  
7 the base to determine the location of the target range list of entries.
- 8 – An offset into the `.debug_rnglists` section ([DW\\_FORM\\_sec\\_offset](#)).  
9 The operand consists of a byte offset from the beginning of the  
10 `.debug_rnglists` section. It is relocatable in a relocatable object file,  
11 and relocated in an executable or shared object file. In the [32-bit  
12 DWARF format](#), this offset is a 4-byte unsigned value; in the [64-bit  
13 DWARF format](#), it is an 8-byte unsigned value (see [Section 7.4 on  
14 page 204](#)).

15 • [rnglistsptr](#)

16 This is an offset into the `.debug_rnglists` section ([DW\\_FORM\\_sec\\_offset](#)).  
17 It consists of a byte offset from the beginning of the `.debug_rnglists`  
18 section. It is relocatable in a relocatable object file, and relocated in an  
19 executable or shared object file. In the [32-bit DWARF format](#), this offset is a  
20 4-byte unsigned value; in the 64-bit DWARF format, it is an 8-byte  
21 unsigned value (see [Section 7.4 on page 204](#)).

22 • [reference](#)

23 There are four types of reference.

- 24 – The first type of reference can identify any debugging information  
25 entry within the containing unit. This type of reference is an offset  
26 from the first byte of the compilation header for the compilation unit  
27 containing the reference. There are five forms for this type of reference.  
28 There are fixed length forms for one, two, four and eight byte offsets  
29 (respectively, [DW\\_FORM\\_ref1](#), [DW\\_FORM\\_ref2](#), [DW\\_FORM\\_ref4](#),  
30 and [DW\\_FORM\\_ref8](#)). There is also an unsigned variable length offset  
31 encoded form that uses unsigned LEB128 numbers  
32 ([DW\\_FORM\\_ref\\_udata](#)). Because this type of reference is within the  
33 containing compilation unit, no relocation of the value is required.
- 34 – The second type of reference can identify any debugging information  
35 entry within a `.debug_info` section; in particular, it may refer to an  
36 entry in a different compilation unit from the unit containing the  
37 reference, and may refer to an entry in a different shared object file.  
38 This type of reference ([DW\\_FORM\\_ref\\_addr](#)) is an offset from the

## Chapter 7. Data Representation

1 beginning of the `.debug_info` section of the target executable or  
2 shared object file, or, for references within a supplementary object file,  
3 an offset from the beginning of the local `.debug_info` section; it is  
4 relocatable in a relocatable object file and frequently relocated in an  
5 executable or shared object file. In the [32-bit DWARF format](#), this  
6 offset is a 4-byte unsigned value; in the [64-bit DWARF format](#), it is an  
7 8-byte unsigned value (see [Section 7.4 on page 204](#)).

8 *A debugging information entry that may be referenced by another compilation*  
9 *unit using `DW_FORM_ref_addr` must have a global symbolic name.*

10 *For a reference from one executable or shared object file to another, the*  
11 *reference is resolved by the debugger to identify the executable or shared object*  
12 *file and the offset into that file's `.debug_info` section in the same fashion as*  
13 *the run time loader, either when the debug information is first read, or when*  
14 *the reference is used.*

- 15 – The third type of reference can identify any debugging information  
16 type entry that has been placed in its own type unit. This type of  
17 reference ([`DW\_FORM\_ref\_sig8`](#)) is the 8-byte type signature (see  
18 [Section 7.32 on page 257](#)) that was computed for the type.
- 19 – The fourth type of reference is a reference from within the `.debug_info`  
20 section of the executable or shared object file to a debugging  
21 information entry in the `.debug_info` section of a supplementary  
22 object file. This type of reference ([`DW\_FORM\_ref\_sup4`](#) or  
23 [`DW\_FORM\_ref\_sup8`](#)) is a 4- or 8-byte offset (respectively) from the  
24 beginning of the `.debug_info` section in the supplementary object file.

25 *The use of compilation unit relative references will reduce the number of*  
26 *link-time relocations and so speed up linking. The use of the second, third and*  
27 *fourth type of reference allows for the sharing of information, such as types,*  
28 *across compilation units, while the fourth type further allows for sharing of*  
29 *information across compilation units from different executables or shared*  
30 *object files.*

31 *A reference to any kind of compilation unit identifies the debugging*  
32 *information entry for that unit, not the preceding header.*

- 33 • [string](#)  
34 A string is a sequence of contiguous non-null bytes followed by one null  
35 byte. A string may be represented:
  - 36 – Immediately in the debugging information entry itself  
37 ([`DW\_FORM\_string`](#)),

## Chapter 7. Data Representation

- As an offset into a string table contained in the `.debug_str` section of the object file (`DW_FORM_strp` or `DW_FORM_strp8`), the `.debug_line_str` section of the object file (`DW_FORM_line_strp`), or as an offset into a string table contained in the `.debug_str` section of a supplementary object file (`DW_FORM_strp_sup` or `DW_FORM_strp_sup8`), `DW_FORM_strp_sup` offsets from the `.debug_info` section of a supplementary object file refer to the local `.debug_str` section of that same file.

In the [32-bit DWARF format](#), the representation of a `DW_FORM_strp`, `DW_FORM_line_strp` or `DW_FORM_strp_sup` value is a 4-byte unsigned offset; in the [64-bit DWARF format](#), it is an 8-byte unsigned offset (see [Section 7.4 on page 204](#)). In both 32-bit and 64-bit formats, the representation of a `DW_FORM_strp8` or `DW_FORM_strp_sup8` value is an 8-byte unsigned offset.

- As an indirect offset into the string table using an index into a table of offsets contained in the `.debug_str_offsets` section of the object file (`DW_FORM_strx`, `DW_FORM_strx1`, `DW_FORM_strx2`, `DW_FORM_strx3` and `DW_FORM_strx4`). The representation of a `DW_FORM_strx` value is an unsigned LEB128 value, which is interpreted as a zero-based index into an array of offsets in the `.debug_str_offsets` section. The representation of a `DW_FORM_strx1`, `DW_FORM_strx2`, `DW_FORM_strx3` or `DW_FORM_strx4` value is a 1-, 2-, 3- or 4-byte unsigned integer value, respectively, which is similarly interpreted. The offset entries in the `.debug_str_offsets` section are described in [Section 7.26 on page 253](#).

Any combination of these three forms may be used within a single compilation.

If the `DW_AT_use_UTF8` attribute is specified for the compilation, partial, skeleton or type unit entry, string values are encoded using the UTF-8 (Unicode Transformation Format-8) from the Universal Character Set standard (ISO/IEC 10646-1:1993). Otherwise, the string representation is unspecified.

*The Unicode Standard Version 3 is fully compatible with ISO/IEC 10646-1:1993. It contains all the same characters and encoding points as ISO/IEC 10646, as well as additional information about the characters and their use.*

## Chapter 7. Data Representation

1 *Earlier versions of DWARF did not specify the representation of strings; for*  
2 *compatibility, this version also does not. However, the UTF-8 representation is*  
3 *strongly recommended.*

- 4 • **stroffsetsptr**

5 This is an offset into the `.debug_str_offsets` section  
6 (`DW_FORM_sec_offset`). It consists of an offset from the beginning of the  
7 `.debug_str_offsets` section to the header of the string offsets information  
8 for the referencing entity. It is relocatable in a relocatable object file, and  
9 relocated in an executable or shared object file. In the **32-bit DWARF**  
10 **format**, this offset is a 4-byte unsigned value; in the **64-bit DWARF format**,  
11 it is an 8-byte unsigned value (see Section 7.4 on page 204).

12 In no case does an attribute use one of the classes `addrptr`, `lineptr`, `loclistsptr`,  
13 `macptr`, `rnglistsptr` or `stroffsetsptr` to point into either the `.debug_info` or  
14 `.debug_str` section.

### 15 7.5.6 Form Encodings

16 The form encodings are listed in Table 7.6 following.

Table 7.6: Attribute form encodings

Form name	Value	Classes
<code>DW_FORM_addr</code>	0x01	<code>address</code>
<i>Reserved</i>	0x02	
<code>DW_FORM_block2</code>	0x03	<code>block</code>
<code>DW_FORM_block4</code>	0x04	<code>block</code>
<code>DW_FORM_data2</code>	0x05	<code>constant</code>
<code>DW_FORM_data4</code>	0x06	<code>constant</code>
<code>DW_FORM_data8</code>	0x07	<code>constant</code>
<code>DW_FORM_string</code>	0x08	<code>string</code>
<code>DW_FORM_block</code>	0x09	<code>block</code>
<code>DW_FORM_block1</code>	0x0a	<code>block</code>
<code>DW_FORM_data1</code>	0x0b	<code>constant</code>
<code>DW_FORM_flag</code>	0x0c	<code>flag</code>
<code>DW_FORM_sdata</code>	0x0d	<code>constant</code>
<code>DW_FORM_strp</code>	0x0e	<code>string</code>

*Continued on next page*

## Chapter 7. Data Representation

Form name	Value	Classes
DW_FORM_udata	0x0f	constant
DW_FORM_ref_addr	0x10	reference
DW_FORM_ref1	0x11	reference
DW_FORM_ref2	0x12	reference
DW_FORM_ref4	0x13	reference
DW_FORM_ref8	0x14	reference
DW_FORM_ref_udata	0x15	reference
DW_FORM_indirect	0x16	(see Section 7.5.3 on page 213)
DW_FORM_sec_offset	0x17	addrptr, lineptr, loclist, loclistsptr, macptr, rnglist, rnglistsptr, stroffsetsptr
DW_FORM_exprloc	0x18	exprval, locdesc
DW_FORM_flag_present	0x19	flag
DW_FORM_strx	0x1a	string
DW_FORM_addrx	0x1b	address
DW_FORM_ref_sup4	0x1c	reference
DW_FORM_strp_sup	0x1d	string
DW_FORM_data16	0x1e	constant
DW_FORM_line_strp	0x1f	string
DW_FORM_ref_sig8	0x20	reference
DW_FORM_implicit_const	0x21	constant
DW_FORM_loclistx	0x22	loclist
DW_FORM_rnglistx	0x23	rnglist
DW_FORM_ref_sup8	0x24	reference
DW_FORM_strx1	0x25	string
DW_FORM_strx2	0x26	string
DW_FORM_strx3	0x27	string
DW_FORM_strx4	0x28	string
DW_FORM_addrx1	0x29	address
DW_FORM_addrx2	0x2a	address
DW_FORM_addrx3	0x2b	address
DW_FORM_addrx4	0x2c	address
DW_FORM_strp8 ‡	0x2d	string

*Continued on next page*

Form name	Value	Classes
DW_FORM_strp_sup8 ‡	0x2e	string

‡ New in DWARF Version 6

## 7.6 Variable Length Data

Integers may be encoded using “Little-Endian Base 128” (LEB128) numbers. LEB128 is a scheme for encoding integers densely that exploits the assumption that most integers are small in magnitude.

*This encoding is equally suitable whether the target machine architecture represents data in big-endian or little-endian byte order. It is “little-endian” only in the sense that it avoids using space to represent the “big” end of an unsigned integer, when the big end is all zeroes or sign extension bits.*

Unsigned LEB128 (ULEB128) numbers are encoded as follows: start at the low order end of an unsigned integer and chop it into 7-bit chunks. Place each chunk into the low order 7 bits of a byte. Typically, several of the high order bytes will be zero, which may be discarded. Emit the remaining bytes in a stream, starting with the low order byte; set the high order bit on each byte except the last emitted byte. The high bit of zero on the last byte indicates to the decoder that it has encountered the last byte.

The integer zero is a special case, consisting of a single zero byte.

Table 7.7 on the following page gives some examples of unsigned LEB128 numbers. The 0x80 in each case is the high order bit of the byte, indicating that an additional byte follows.

The encoding for signed, two’s complement LEB128 (SLEB128) numbers is similar, except that the criterion for discarding high order bytes is not whether they are zero, but whether they consist entirely of sign extension bits. Consider the 4-byte integer -2. The three high level bytes of the number are sign extension, thus LEB128 would represent it as a single byte containing the low order 7 bits, with the high order bit cleared to indicate the end of the byte stream. Note that there is nothing within the LEB128 representation that indicates whether an encoded number is signed or unsigned. The decoder must know what type of number to expect. Table 7.7 on the next page gives some examples of unsigned LEB128 numbers and Table 7.8 on the following page gives some examples of signed LEB128 numbers.

## Chapter 7. Data Representation

- 1 *Some producers may choose to insert padding or alignment bytes by retaining (not*  
2 *discarding) one or more high-order bytes that would not affect the decoded value.*  
3 *Appendix C on page 294 gives algorithms for encoding and decoding these forms.*

Table 7.7: Examples of unsigned LEB128 encodings

Number	First byte	Second byte
2	2	—
127	127	—
128	0 + 0x80	1
129	1 + 0x80	1
12857	57 + 0x80	100

Table 7.8: Examples of signed LEB128 encodings

Number	First byte	Second byte
2	2	—
-2	0x7e	—
127	127 + 0x80	0
-127	1 + 0x80	0x7f
128	0 + 0x80	1
-128	0 + 0x80	0x7f
129	1 + 0x80	1
-129	0x7f + 0x80	0x7e

## 4 **7.7 DWARF Expressions and Location Descriptions**

### 5 **7.7.1 DWARF Expressions**

- 6 A DWARF expression is stored in a block of contiguous bytes. The bytes form a  
7 sequence of operations. Each operation is a 1-byte code that identifies that  
8 operation, followed by zero or more bytes of additional data. The encodings for  
9 the operations are described in Table 7.9 on the following page.

## Chapter 7. Data Representation

Table 7.9: DWARF operation encodings

Operation	Code	No. of Operands	Notes
<i>Reserved</i>	0x01	-	
<i>Reserved</i>	0x02	-	
DW_OP_addr	0x03	1	constant address (size is target specific)
<i>Reserved</i>	0x04	-	
<i>Reserved</i>	0x05	-	
DW_OP_deref	0x06	0	
<i>Reserved</i>	0x07	-	
DW_OP_const1u	0x08	1	1-byte constant
DW_OP_const1s	0x09	1	1-byte constant
DW_OP_const2u	0x0a	1	2-byte constant
DW_OP_const2s	0x0b	1	2-byte constant
DW_OP_const4u	0x0c	1	4-byte constant
DW_OP_const4s	0x0d	1	4-byte constant
DW_OP_const8u	0x0e	1	8-byte constant
DW_OP_const8s	0x0f	1	8-byte constant
DW_OP_constu	0x10	1	ULEB128 constant
DW_OP_consts	0x11	1	SLEB128 constant
DW_OP_dup	0x12	0	
DW_OP_drop	0x13	0	
DW_OP_over	0x14	0	
DW_OP_pick	0x15	1	1-byte stack index
DW_OP_swap	0x16	0	
DW_OP_rot	0x17	0	
DW_OP_xderef	0x18	0	
DW_OP_abs	0x19	0	
DW_OP_and	0x1a	0	
DW_OP_div	0x1b	0	
DW_OP_minus	0x1c	0	

*Continued on next page*



## Chapter 7. Data Representation

Operation	Code	No. of Operands	Notes
DW_OP_mod	0x1d	0	
DW_OP_mul	0x1e	0	
DW_OP_neg	0x1f	0	
DW_OP_not	0x20	0	
DW_OP_or	0x21	0	
DW_OP_plus	0x22	0	
DW_OP_plus_uconst	0x23	1	ULEB128 addend
DW_OP_shl	0x24	0	
DW_OP_shr	0x25	0	
DW_OP_shra	0x26	0	
DW_OP_xor	0x27	0	
DW_OP_bra	0x28	1	signed 2-byte constant
DW_OP_eq	0x29	0	
DW_OP_ge	0x2a	0	
DW_OP_gt	0x2b	0	
DW_OP_le	0x2c	0	
DW_OP_lt	0x2d	0	
DW_OP_ne	0x2e	0	
DW_OP_skip	0x2f	1	signed 2-byte constant
DW_OP_lit0	0x30	0	
DW_OP_lit1	0x31	0	
...			
DW_OP_lit31	0x4f	0	literals 0 .. 31 = (DW_OP_lit0 + literal)
DW_OP_reg0	0x50	0	
DW_OP_reg1	0x51	0	
...			
DW_OP_reg31	0x6f	0	reg 0 .. 31 = (DW_OP_reg0 + regnum)
DW_OP_breg0	0x70	1	SLEB128 offset
DW_OP_breg1	0x71	1	
...			
DW_OP_breg31	0x8f	1	base register 0 .. 31 = (DW_OP_breg0 + regnum)

*Continued on next page*

## Chapter 7. Data Representation

Operation	Code	No. of Operands	Notes
DW_OP_regx	0x90	1	ULEB128 register
DW_OP_fbreg	0x91	1	SLEB128 offset
DW_OP_bregx	0x92	2	ULEB128 register, SLEB128 offset
DW_OP_piece	0x93	1	ULEB128 size of piece
DW_OP_deref_size	0x94	1	1-byte size of data retrieved
DW_OP_xderef_size	0x95	1	1-byte size of data retrieved
DW_OP_nop	0x96	0	
DW_OP_push_object_address	0x97	0	
DW_OP_call2	0x98	1	2-byte offset of DIE
DW_OP_call4	0x99	1	4-byte offset of DIE
DW_OP_call_ref	0x9a	1	4- or 8-byte offset of DIE
DW_OP_form_tls_address	0x9b	0	
DW_OP_call_frame_cfa	0x9c	0	
DW_OP_bit_piece	0x9d	2	ULEB128 size, ULEB128 offset
DW_OP_implicit_value	0x9e	2	ULEB128 size, block of that size
DW_OP_stack_value	0x9f	0	
DW_OP_implicit_pointer	0xa0		4- or 8-byte offset of DIE, SLEB128 constant offset
DW_OP_addrx	0xa1	1	ULEB128 indirect address
DW_OP_constx	0xa2	1	ULEB128 indirect constant
DW_OP_entry_value	0xa3	2	ULEB128 size, block of that size
DW_OP_const_type	0xa4	3	ULEB128 type entry offset, 1-byte size, constant value
DW_OP_regval_type	0xa5	2	ULEB128 register number, ULEB128 constant offset

*Continued on next page*

Operation	Code	No. of Operands	Notes
DW_OP_deref_type	0xa6	2	1-byte size, ULEB128 type entry offset
DW_OP_xderef_type	0xa7	2	1-byte size, ULEB128 type entry offset
DW_OP_convert	0xa8	1	ULEB128 type entry offset
DW_OP_reinterpret	0xa9	1	ULEB128 type entry offset
DW_OP_regval_bits ‡	0xaa	1 TOS † TOS - 1	ULEB128 field size, integer bit offset, integer register number
DW_OP_user_extended ‡	0xdf	1 +	
DW_OP_lo_user	0xe0		
DW_OP_hi_user	0xff		

‡ New in DWARF Version 6

† TOS indicates parameter on top of stack

## 7.7.2 Location Descriptions

A location description is used to compute the location of a variable or other entity.

## 7.7.3 Location Lists

Each entry in a location list is either a location list entry, a base address entry, a default location entry or an end-of-list entry.

Each entry begins with an unsigned 1-byte code that indicates the kind of entry that follows. The encodings for these constants are given in Table 7.10 on the next page.

Table 7.10: Location list entry encoding values

Location list entry encoding name	Value
DW_LLE_end_of_list	0x00
DW_LLE_base_addressx	0x01
DW_LLE_startx_endx	0x02
DW_LLE_startx_length	0x03
DW_LLE_offset_pair	0x04
DW_LLE_default_location	0x05
DW_LLE_base_address	0x06
DW_LLE_start_end	0x07
DW_LLE_start_length	0x08
DW_LLE_include_loclist ‡	0x09
DW_LLE_include_loclistx ‡	0x0a
DW_LLE_lo_user ‡	0xc0
DW_LLE_hi_user ‡	0xff

‡ New in DWARF Version 6

1 If a vendor defines a vendor-specific kind of location list entry, the kind code  
 2 must be immediately followed by an unsigned LEB128 value that specifies the  
 3 length of all remaining bytes (not including either the kind or the length itself)  
 4 for that entry.

## 5 7.8 Base Type Attribute Encodings

6 The encodings of the constants used in the `DW_AT_encoding` attribute are given  
 7 in Table 7.11.

Table 7.11: Base type encoding values

Base type encoding name	Value
DW_ATE_address	0x01
DW_ATE_boolean	0x02
DW_ATE_complex_float	0x03
DW_ATE_float	0x04
<i>Continued on next page</i>	

## Chapter 7. Data Representation

Base type encoding name	Value
DW_ATE_signed	0x05
DW_ATE_signed_char	0x06
DW_ATE_unsigned	0x07
DW_ATE_unsigned_char	0x08
DW_ATE_imaginary_float	0x09
DW_ATE_packed_decimal	0x0a
DW_ATE_numeric_string	0x0b
DW_ATE_edited	0x0c
DW_ATE_signed_fixed	0x0d
DW_ATE_unsigned_fixed	0x0e
DW_ATE_decimal_float	0x0f
DW_ATE_UTF	0x10
DW_ATE_UCS	0x11
DW_ATE_ASCII	0x12
DW_ATE_complex_signed ‡	0x13
DW_ATE_imaginary_signed ‡	0x14
DW_ATE_complex_unsigned ‡	0x15
DW_ATE_imaginary_unsigned ‡	0x16
DW_ATE_signed_bitint ‡	0x17
DW_ATE_unsigned_bitint ‡	0x18
DW_ATE_lo_user	0x80
DW_ATE_hi_user	0xff

‡ New in DWARF Version 6

- 1 The encodings of the constants used in the `DW_AT_decimal_sign` attribute are  
 2 given in Table 7.12 on the next page.

Table 7.12: Decimal sign encodings

Decimal sign code name	Value
DW_DS_unsigned	0x01
DW_DS_leading_overpunch	0x02
DW_DS_trailing_overpunch	0x03
DW_DS_leading_separate	0x04
DW_DS_trailing_separate	0x05

1 The encodings of the constants used in the `DW_AT_endianity` attribute are given  
2 in Table 7.13.

Table 7.13: Endianity encodings

Endian code name	Value
DW_END_default	0x00
DW_END_big	0x01
DW_END_little	0x02
DW_END_lo_user	0x40
DW_END_hi_user	0xff

## 3 7.9 Accessibility Codes

4 The encodings of the constants used in the `DW_AT_accessibility` attribute are  
5 given in Table 7.14.

Table 7.14: Accessibility encodings

Accessibility code name	Value
DW_ACCESS_public	0x01
DW_ACCESS_protected	0x02
DW_ACCESS_private	0x03

## 6 7.10 Visibility Codes

7 The encodings of the constants used in the `DW_AT_visibility` attribute are given  
8 in Table 7.15 on the following page.

Table 7.15: Visibility encodings

Visibility code name	Value
<a href="#">DW_VIS_local</a>	0x01
<a href="#">DW_VIS_exported</a>	0x02
<a href="#">DW_VIS_qualified</a>	0x03

## 7.11 Virtuality Codes

The encodings of the constants used in the [DW\\_AT\\_virtuality](#) attribute are given in Table 7.16.

Table 7.16: Virtuality encodings

Virtuality code name	Value
<a href="#">DW_VIRTUALITY_none</a>	0x00
<a href="#">DW_VIRTUALITY_virtual</a>	0x01
<a href="#">DW_VIRTUALITY_pure_virtual</a>	0x02

The value [DW\\_VIRTUALITY\\_none](#) is equivalent to the absence of the [DW\\_AT\\_virtuality](#) attribute.

## 7.12 Source Languages

The encodings of the constants used in the [DW\\_AT\\_language\\_name](#) attribute are given in Table 7.17 on the following page. Table 7.17 on the next page also shows the default lower bound, if any, assumed for an omitted [DW\\_AT\\_lower\\_bound](#) attribute in the context of a [DW\\_TAG\\_subrange\\_type](#) debugging information entry for each defined language.

---

*Continued on next page*

## Chapter 7. Data Representation

Language name	Value	Default Lower Bound
---------------	-------	---------------------

Table 7.17: Language encodings

Language name	Value	Default Lower Bound
DW_LNAME_Ada	0x0001	1
DW_LNAME_BLISS	0x0002	0
DW_LNAME_C	0x0003	0
DW_LNAME_C_plus_plus	0x0004	0
DW_LNAME_Cobol	0x0005	1
DW_LNAME_Crystal ‡	0x0006	0
DW_LNAME_D	0x0007	0
DW_LNAME_Dylan	0x0008	0
DW_LNAME_Fortran	0x0009	1
DW_LNAME_Go	0x000a	0
DW_LNAME_Haskell	0x000b	0
DW_LNAME_Java	0x000c	0
DW_LNAME_Julia	0x000d	1
DW_LNAME_Kotlin ‡	0x000e	0
DW_LNAME_Modula2	0x000f	1
DW_LNAME_Modula3	0x0010	1
DW_LNAME_ObjC	0x0011	0
DW_LNAME_ObjC_plus_plus	0x0012	0
DW_LNAME_OCaml	0x0013	0
DW_LNAME_OpenCL_C <sup>7</sup>	0x0014	0
DW_LNAME_Pascal	0x0015	1
DW_LNAME_PLI	0x0016	1
DW_LNAME_Python	0x0017	0
DW_LNAME_RenderScript	0x0018	0
DW_LNAME_Rust	0x0019	0
DW_LNAME_Swift	0x001a	0
DW_LNAME_UPC	0x001b	0
DW_LNAME_Zig ‡	0x001c	0

*Continued on next page*

<sup>7</sup>Formerly DW\_LANG\_OpenCL in DWARF Version 5.



Language name	Value	Default Lower Bound
<a href="#">DW_LNAME_Assembly</a> ‡	0x001d	0
<a href="#">DW_LNAME_C_sharp</a> ‡	0x001e	0
<a href="#">DW_LNAME_Mojo</a> ‡	0x001f	0
<a href="#">DW_LNAME_GLSL</a> ‡	0x0020	0
<a href="#">DW_LNAME_GLSL_ES</a> ‡	0x0021	0
<a href="#">DW_LNAME_HLSL</a> ‡	0x0022	0
<a href="#">DW_LNAME_OpenCL_CPP</a> ‡	0x0023	0
<a href="#">DW_LNAME_CPP_for_OpenCL</a> ‡	0x0024	0
<a href="#">DW_LNAME_SYCL</a> ‡	0x0025	0
<a href="#">DW_LNAME_Ruby</a> ‡	0x0026	0
<a href="#">DW_LNAME_lo_user</a>	0x8000	
<a href="#">DW_LNAME_hi_user</a>	0xffff	

‡ Base language is new in DWARF Version 6

## 7.13 Address Class Encodings

The value of the common address class encoding [DW\\_ADDR\\_none](#) is 0.

## 7.14 Identifier Case

The encodings of the constants used in the [DW\\_AT\\_identifier\\_case](#) attribute are given in Table 7.18.

Table 7.18: Identifier case encodings

Identifier case name	Value
<a href="#">DW_ID_case_sensitive</a>	0x00
<a href="#">DW_ID_up_case</a>	0x01
<a href="#">DW_ID_down_case</a>	0x02
<a href="#">DW_ID_case_insensitive</a>	0x03

## 7.15 Calling Convention Encodings

The encodings of the constants used in the [DW\\_AT\\_calling\\_convention](#) attribute are given in Table 7.19.

Table 7.19: Calling convention encodings

Calling convention name	Value
<a href="#">DW_CC_normal</a>	0x01
<a href="#">DW_CC_program</a>	0x02
<a href="#">DW_CC_nocall</a>	0x03
<a href="#">DW_CC_pass_by_reference</a>	0x04
<a href="#">DW_CC_pass_by_value</a>	0x05
<a href="#">DW_CC_lo_user</a>	0x40
<a href="#">DW_CC_hi_user</a>	0xff

## 7.16 Inline Codes

The encodings of the constants used in the [DW\\_AT\\_inline](#) attribute are given in Table 7.20.

Table 7.20: Inline encodings

Inline code name	Value
<a href="#">DW_INL_not_inlined</a>	0x00
<a href="#">DW_INL_inlined</a>	0x01
<a href="#">DW_INL_declared_not_inlined</a>	0x02
<a href="#">DW_INL_declared_inlined</a>	0x03

## 7.17 Array Ordering

The encodings of the constants used in the [DW\\_AT\\_ordering](#) attribute are given in Table 7.21 on the next page.

Table 7.21: Ordering encodings

Ordering name	Value
DW_ORD_row_major	0x00
DW_ORD_col_major	0x01

## 7.18 Discriminant Lists

The descriptors used in the `DW_AT_discr_list` attribute are encoded as 1-byte constants. The defined values are given in Table 7.22.

Table 7.22: Discriminant descriptor encodings

Descriptor name	Value
DW_DSC_label	0x00
DW_DSC_range	0x01

## 7.19 Name Index Table

The version number in the name index table header is 6.

The name index attributes and their encodings are listed in Table 7.23.

Table 7.23: Name index attribute encodings

Attribute name	Value	Form/Class
DW_IDX_compile_unit	1	constant
DW_IDX_type_unit	2	constant
DW_IDX_die_offset	3	reference
DW_IDX_parent	4	constant
DW_IDX_type_hash	5	DW_FORM_data8
DW_IDX_external ‡	6	flag
DW_IDX_lo_user	0x2000	
DW_IDX_hi_user	0x3fff	

‡ New in DWARF Version 6

1 *It is suggested that producers should use the form code `DW_FORM_flag_present` for the*  
 2 *`DW_IDX_external` attribute for abbreviation codes that represent external names.*

3 The abbreviations table ends with an entry consisting of a single 0 byte for the  
 4 abbreviation code. The size of the table given by `abbrev_table_size` may  
 5 include optional padding following the terminating 0 byte.

## 6 7.20 Defaulted Member Encodings

7 The encodings of the constants used in the `DW_AT_defaulted` attribute are given  
 8 in Table 7.24 following.

Table 7.24: Defaulted attribute encodings

Defaulted name	Value
<code>DW_DEFAULTED_no</code>	0x00
<code>DW_DEFAULTED_in_class</code>	0x01
<code>DW_DEFAULTED_out_of_class</code>	0x02

## 9 7.21 Address Range Table

10 Each `.debug_aranges` section contribution begins with a header containing:

- 11 1. `unit_length` ([initial length](#))  
 12 A 4-byte or 12-byte length containing the length of the set of entries for this  
 13 compilation unit, not including the length field itself (see [Section 7.4 on](#)  
 14 [page 204](#)).
- 15 2. `version` ([uhalf](#))  
 16 A 2-byte version identifier representing the version of the DWARF  
 17 information for the address range table.  
 18 This value in this field is 2.
- 19 3. `debug_info_offset` ([section offset](#))  
 20 A 4-byte or 8-byte offset into the `.debug_info` section of the compilation unit  
 21 header. In the [32-bit DWARF format](#), this is a 4-byte unsigned offset; in the  
 22 [64-bit DWARF format](#), this is an 8-byte unsigned offset (see [Section 7.4 on](#)  
 23 [page 204](#)).
- 24 4. `address_size` ([ubyte](#))  
 25 A 1-byte unsigned integer containing the size in bytes of an address.

1 5. *reserved*<sup>8</sup> (ubyte, MBZ)

2 This header is followed by a series of tuples. Each tuple consists of an address  
3 and a length. The address and length size are each given by the `address_size`  
4 field of the header. The first tuple following the header in each set begins at an  
5 offset that is a multiple of the size of a single tuple (that is, twice the size of an  
6 address). The header is padded, if necessary, to that boundary. Each set of tuples  
7 is terminated by a 0 for the address and a 0 for the length.

## 8 7.22 Line Number Information

9 The version number in the line number program header is 6.

10 The boolean values “true” and “false” used by the line number information  
11 program are encoded as a single byte containing the value 0 for “false,” and a  
12 non-zero value for “true.”

13 The encodings for the standard opcodes are given in Table 7.25.

Table 7.25: Line number standard opcode encodings

Opcode name	Value
DW_LNS_extended_op ‡	0x00
DW_LNS_copy	0x01
DW_LNS_advance_pc	0x02
DW_LNS_advance_line	0x03
DW_LNS_set_file	0x04
DW_LNS_set_column	0x05
DW_LNS_negate_stmt	0x06
DW_LNS_set_basic_block	0x07
DW_LNS_const_add_pc	0x08
DW_LNS_fixed_advance_pc	0x09
DW_LNS_set_prologue_end	0x0a
DW_LNS_set_epilogue_begin	0x0b
DW_LNS_set_isa	0x0c

‡ New in DWARF Version 6

---

<sup>8</sup>This allows backward compatible support of the deprecated `segment_selector_size` field which was defined in DWARF Version 5 and earlier.

## Chapter 7. Data Representation

1 The encodings for the extended opcodes are given in Table 7.26.

Table 7.26: Line number extended opcode encodings

Opcode name	Value
DW_LNE_end_sequence	0x01
DW_LNE_set_address	0x02
<i>Reserved</i>	0x03 <sup>9</sup>
DW_LNE_set_discriminator	0x04
DW_LNE_padding ‡	0x05
DW_LNE_lo_user	0x80
DW_LNE_hi_user	0xff

‡ New in DWARF Version 6

2 The encodings for the line number header entry formats are given in Table 7.27.

Table 7.27: Line number header entry format encodings

Line number header entry format name	Value
DW_LNCT_path	0x1
DW_LNCT_directory_index	0x2
DW_LNCT_timestamp	0x3
DW_LNCT_size	0x4
DW_LNCT_MD5	0x5
DW_LNCT_source ‡	0x6
DW_LNCT_URL ‡	0x7
DW_LNCT_lo_user	0x2000
DW_LNCT_hi_user	0x3fff

‡ New in DWARF Version 6

---

<sup>9</sup>Code 0x03 is reserved to allow backward compatible support of the DW\_LNE\_define\_file operation which was defined in DWARF Version 4 and earlier.

1 **7.23 Macro Information**

2 The version number in the macro information header is 5.

3 The source line numbers and source file indices encoded in the macro  
4 information section are represented as unsigned LEB128 numbers.

5 The macro information entry type is encoded as a single unsigned byte. The  
6 encodings are given in [Table 7.28 on the next page](#).

## Chapter 7. Data Representation

Table 7.28: Macro information entry type encodings

Macro information entry type name	Value
DW_MACRO_define	0x01
DW_MACRO_undef	0x02
DW_MACRO_start_file	0x03
DW_MACRO_end_file	0x04
DW_MACRO_define_strp	0x05
DW_MACRO_undef_strp	0x06
DW_MACRO_import	0x07
<i>Reserved</i>	0x08 <sup>10</sup>
<i>Reserved</i>	0x09 <sup>11</sup>
<i>Reserved</i>	0x0a <sup>12</sup>
DW_MACRO_define_strx	0x0b
DW_MACRO_undef_strx	0x0c
DW_MACRO_define_sup4 ‡	0x0d
DW_MACRO_define_sup8 ‡	0x0e
DW_MACRO_undef_sup4 ‡	0x0f
DW_MACRO_undef_sup8 ‡	0x10
DW_MACRO_import_sup4 ‡	0x11
DW_MACRO_import_sup8 ‡	0x12
DW_MACRO_lo_user	0xe0
DW_MACRO_hi_user	0xff

‡ New in DWARF Version 6

<sup>10</sup>Code 0x08 is reserved to allow backward compatible support of the DW\_MACRO\_define\_sup entry type that was defined (only) in DWARF Version 5.

<sup>11</sup>Code 0x09 is reserved to allow backward compatible support of the DW\_MACRO\_undef\_sup entry type that was defined (only) in DWARF Version 5.

<sup>12</sup>Code 0x0a is reserved to allow backward compatible support of the DW\_MACRO\_import\_sup entry type that was defined (only) in DWARF Version 5.



## 7.24 Call Frame Information

In the **32-bit DWARF format**, the value of the CIE id in the CIE header is 0xffffffff; in the **64-bit DWARF format**, the value is 0xffffffffffffffff.

The value of the CIE version number is 4.

Call frame instructions are encoded in one or more bytes. The primary opcode is encoded in the high order two bits of the first byte (that is, opcode = byte  $\gg$  6).

An operand or extended opcode may be encoded in the low order 6 bits.

Additional operands are encoded in subsequent bytes. The instructions and their encodings are presented in Table 7.29.

Table 7.29: Call frame instruction encodings

Instruction	High 2 Bits	Low 6 Bits	Operand 1	Operand 2
DW_CFA_advance_loc	0x1	delta		
DW_CFA_offset	0x2	register	ULEB128 offset	
DW_CFA_restore	0x3	register		
DW_CFA_nop	0	0		
DW_CFA_set_loc	0	0x01	address	
DW_CFA_advance_loc1	0	0x02	1-byte delta	
DW_CFA_advance_loc2	0	0x03	2-byte delta	
DW_CFA_advance_loc4	0	0x04	4-byte delta	
DW_CFA_offset_extended	0	0x05	ULEB128 register	ULEB128 offset
DW_CFA_restore_extended	0	0x06	ULEB128 register	
DW_CFA_undefined	0	0x07	ULEB128 register	
DW_CFA_same_value	0	0x08	ULEB128 register	
DW_CFA_register	0	0x09	ULEB128 register	ULEB128 offset
DW_CFA_remember_state	0	0x0a		
DW_CFA_restore_state	0	0x0b		
DW_CFA_def_cfa	0	0x0c	ULEB128 register	ULEB128 offset
DW_CFA_def_cfa_register	0	0x0d	ULEB128 register	
DW_CFA_def_cfa_offset	0	0x0e	ULEB128 offset	
DW_CFA_def_cfa_expression	0	0x0f	exprloc	
DW_CFA_expression	0	0x10	ULEB128 register	exprloc

*Continued on next page*

Instruction	High 2 Bits	Low 6 Bits	Operand 1	Operand 2
<a href="#">DW_CFA_offset_extended_sf</a>	0	0x11	ULEB128 register	SLEB128 offset
<a href="#">DW_CFA_def_cfa_sf</a>	0	0x12	ULEB128 register	SLEB128 offset
<a href="#">DW_CFA_def_cfa_offset_sf</a>	0	0x13	SLEB128 offset	
<a href="#">DW_CFA_val_offset</a>	0	0x14	ULEB128 register	ULEB128 offset
<a href="#">DW_CFA_val_offset_sf</a>	0	0x15	ULEB128 register	SLEB128 offset
<a href="#">DW_CFA_val_expression</a>	0	0x16	ULEB128 register	exprloc
<a href="#">DW_CFA_lo_user</a>	0	0x1c		
<a href="#">DW_CFA_hi_user</a>	0	0x3f		

## 7.25 Range List Entries for Non-contiguous Address Ranges

Each entry in a range list (see Section 2.17.3 on page 53) is either a range list entry, a base address selection entry, or an end-of-list entry.

Each entry begins with an unsigned 1-byte code that indicates the kind of entry that follows. The encodings for these constants are given in Table 7.30.

Table 7.30: Range list entry encoding values

Range list entry encoding name	Value
<a href="#">DW_RLE_end_of_list</a>	0x00
<a href="#">DW_RLE_base_addressx</a>	0x01
<a href="#">DW_RLE_startx_endx</a>	0x02
<a href="#">DW_RLE_startx_length</a>	0x03
<a href="#">DW_RLE_offset_pair</a>	0x04
<a href="#">DW_RLE_base_address</a>	0x05
<a href="#">DW_RLE_start_end</a>	0x06
<a href="#">DW_RLE_start_length</a>	0x07
<a href="#">DW_RLE_include_rnglist ‡</a>	0x08
<a href="#">DW_RLE_include_rnglistx ‡</a>	0x09
<a href="#">DW_RLE_lo_user ‡</a>	0xc0
<a href="#">DW_RLE_hi_user ‡</a>	0xff

‡ New in DWARF Version 6

1 If a vendor defines a vendor-specific kind of range list entry, the kind code must  
2 be immediately followed by an unsigned LEB128 value that specifies the length  
3 of all remaining bytes (not including either the kind or the length itself) for that  
4 entry.

5 For a range list to be specified, the base address of the corresponding compilation  
6 unit must be defined (see Section 3.1.1 on page 62).

## 7 7.26 String Offsets Table

8 Each `.debug_str_offsets` or `.debug_str_offsets.dwo` section contribution  
9 begins with a header containing:

- 10 1. `unit_length` (initial length)  
11 A 4-byte or 12-byte length containing the length of the set of entries for this  
12 compilation unit, not including the length field itself (see Section 7.4 on  
13 page 204).

14 The DWARF format used for the string offsets table is not required to match  
15 the format used by other sections describing the same compilation unit.

- 16 2. `version` (uhalf)  
17 A 2-byte version identifier containing the value 5.
- 18 3. `padding` (uhalf)  
19 Reserved to DWARF (must be zero).

20 This header is followed by a series of string table offset entries that have the same  
21 representation as `DW_FORM_strp`. For the 32-bit DWARF format, each offset is 4  
22 bytes long; for the 64-bit DWARF format, each offset is 8 bytes long.

23 The `DW_AT_str_offsets` attribute points to the header. The entries following the  
24 header are indexed sequentially, starting from 0.

## 25 7.27 Address Table

26 Each `.debug_addr` section contribution begins with a header containing:

- 27 1. `unit_length` (initial length)  
28 A 4-byte or 12-byte length containing the length of the set of entries for this  
29 compilation unit, not including the length field itself (see Section 7.4 on  
30 page 204).

1 2. `version` (uhalf)

2 A 2-byte version identifier containing the value 5.

3 3. `address_size` (ubyte)

4 A 1-byte unsigned integer containing the size in bytes of an address on the  
5 target system.

6 4. *reserved*<sup>13</sup> (ubyte, MBZ)

7 This header is followed by a series of addresses where the address size is given  
8 by the `address_size` field of the header.

9 The `DW_AT_addr_base` attribute points to the first entry following the header.  
10 The entries are indexed sequentially from this base entry, starting from 0.

## 11 7.28 Range List Table

12 Each `.debug_rnglists` and `.debug_rnglists.dwo` section contribution begins  
13 with a header containing:

14 1. `unit_length` (initial length)

15 A 4-byte or 12-byte length containing the length of the set of entries for this  
16 compilation unit, not including the length field itself (see Section 7.4 on  
17 page 204).

18 2. `version` (uhalf)

19 A 2-byte version identifier containing the value 5.

20 3. `address_size` (ubyte)

21 A 1-byte unsigned integer containing the size in bytes of an address on the  
22 target system.

23 4. *reserved*<sup>14</sup> (ubyte, MBZ)

24 5. `offset_entry_count` (uword)

25 A 4-byte count of the number of offsets that follow the header. This count  
26 may be zero.

27 Immediately following the header is an array of offsets. This array is followed by  
28 a series of range lists.

---

<sup>13</sup>This allows backward compatible support of the deprecated `segment_selector_size` field which was defined in DWARF Version 5 and earlier.

<sup>14</sup>This allows backward compatible support of the deprecated `segment_selector_size` field which was defined in DWARF Version 5 and earlier.

1 If the `offset_entry_count` is non-zero, there is one offset for each range list. The  
2 contents of the  $i^{\text{th}}$  offset is the offset (an unsigned integer) from the beginning of  
3 the offset array to the location of the  $i^{\text{th}}$  range list. In the 32-bit DWARF format,  
4 each offset is 4-bytes in size; in the 64-bit DWARF format, each offset is 8-bytes in  
5 size (see Section 7.4 on page 204).

6 *If the `offset_entry_count` is zero, then `DW_FORM_rnglistx` cannot be used to access  
7 a range list; `DW_FORM_sec_offset` must be used instead. If the `offset_entry_count`  
8 is non-zero, then `DW_FORM_rnglistx` may be used to access a range list.*

9 Range lists are described in Section 2.17.3 on page 53.

10 The `DW_AT_rnglists_base` attribute points to the first offset following the header.  
11 The range lists are referenced by the index of the position of their corresponding  
12 offset in the array of offsets, which indirectly specifies the offset to the target list.

## 13 7.29 Location List Table

14 Each `.debug_loclists` or `.debug_loclists.dwo` section contribution begins with  
15 a header containing:

- 16 1. `unit_length` (initial length)  
17 A 4-byte or 12-byte length containing the length of the set of entries for this  
18 compilation unit, not including the length field itself (see Section 7.4 on  
19 page 204).
- 20 2. `version` (uhalf)  
21 A 2-byte version identifier containing the value 5.
- 22 3. `address_size` (ubyte)  
23 A 1-byte unsigned integer containing the size in bytes of an address on the  
24 target system.
- 25 4. `reserved`<sup>15</sup> (ubyte, MBZ)
- 26 5. `offset_entry_count` (uword)  
27 A 4-byte count of the number of offsets that follow the header. This count  
28 may be zero.

29 Immediately following the header is an array of offsets. This array is followed by  
30 a series of location lists.

31 If the `offset_entry_count` is non-zero, there is one offset for each location list.  
32 The contents of the  $i^{\text{th}}$  offset is the offset (an unsigned integer) from the

---

<sup>15</sup>This allows backward compatible support of the deprecated `segment_selector_size` field which was defined in DWARF Version 5 and earlier.

1 beginning of the offset array to the location of the  $i^{\text{th}}$  location list. In the 32-bit  
 2 DWARF format, each offset is 4-bytes in size; in the 64-bit DWARF format, each  
 3 offset is 8-bytes in size (see Section 7.4 on page 204).

4 *If the `offset_entry_count` is zero, then `DW_FORM_loclistx` cannot be used to access a*  
 5 *location list; `DW_FORM_sec_offset` must be used instead. If the `offset_entry_count`*  
 6 *is non-zero, then `DW_FORM_loclistx` may be used to access a location list.*

7 Location lists are described in Section 2.6.2 on page 43.

8 The `DW_AT_loclists_base` attribute points to the first offset following the header.  
 9 The location lists are referenced by the index of the position of their  
 10 corresponding offset in the array of offsets, which indirectly specifies the offset to  
 11 the target list.

## 12 7.30 Dependencies and Constraints

13 The debugging information in this format is intended to exist in sections of an  
 14 object file, or an equivalent separate file or database, having names beginning  
 15 with the prefix ".debug\_" (see Appendix G on page 436 for a complete list of such  
 16 names). Except as specifically specified, this information is not aligned on 2-, 4-  
 17 or 8-byte boundaries. Consequently:

- 18 • For the **32-bit DWARF format** and a target architecture with 32-bit  
 19 addresses, an assembler or compiler must provide a way to produce 2-byte  
 20 and 4-byte quantities without alignment restrictions, and the linker must be  
 21 able to relocate a 4-byte address or section offset that occurs at an arbitrary  
 22 alignment.
- 23 • For the **32-bit DWARF format** and a target architecture with 64-bit  
 24 addresses, an assembler or compiler must provide a way to produce 2-byte,  
 25 4-byte and 8-byte quantities without alignment restrictions, and the linker  
 26 must be able to relocate an 8-byte address or 4-byte section offset that  
 27 occurs at an arbitrary alignment.
- 28 • For the **64-bit DWARF format** and a target architecture with 32-bit  
 29 addresses, an assembler or compiler must provide a way to produce 2-byte,  
 30 4-byte and 8-byte quantities without alignment restrictions, and the linker  
 31 must be able to relocate a 4-byte address or 8-byte section offset that occurs  
 32 at an arbitrary alignment.

33 *It is expected that this will be required only for very large 32-bit programs or by*  
 34 *those architectures which support a mix of 32-bit and 64-bit code and data within*  
 35 *the same executable object.*

- For the [64-bit DWARF format](#) and a target architecture with 64-bit addresses, an assembler or compiler must provide a way to produce 2-byte, 4-byte and 8-byte quantities without alignment restrictions, and the linker must be able to relocate an 8-byte address or section offset that occurs at an arbitrary alignment.

## 7.31 Integer Representation Names

The sizes of the integers used in the lookup by name, lookup by address, line number, call frame information and other sections are given in [Table 7.31](#).

Table 7.31: Integer representation names

Representation name	Representation
sbyte	signed, 1-byte integer
ubyte	unsigned, 1-byte integer
uhalf	unsigned, 2-byte integer
uword	unsigned, 4-byte integer

## 7.32 Type Signature Computation

A type signature is used by a DWARF consumer to resolve type references to the type definitions that are contained in type units (see [Section 3.1.4 on page 71](#)).

*A type signature is computed only by a DWARF producer; a consumer need only compare two type signatures to check for equality.*

The type signature for a type  $T_0$  is formed from the MD5<sup>16</sup> digest of a flattened description of the type. The flattened description of the type is a byte sequence derived from the DWARF encoding of the type as follows:

1. Start with an empty sequence  $S$  and a list  $V$  of visited types, where  $V$  is initialized to a list containing the type  $T_0$  as its single element. Elements in  $V$  are indexed from 1, so that  $V[1]$  is  $T_0$ .

<sup>16</sup>MD5 Message Digest Algorithm, R.L. Rivest, RFC 1321, April 1992

## Chapter 7. Data Representation

- 1       2. If the debugging information entry represents a type that is nested inside  
2       another type or a namespace, append to S the type's context as follows: For  
3       each surrounding type or namespace, beginning with the outermost such  
4       construct, append the letter 'C', the DWARF tag of the construct, and the  
5       name (taken from the [DW\\_AT\\_name](#) attribute) of the type or namespace  
6       (including its trailing null byte).
- 7       3. Append to S the letter 'D', followed by the DWARF tag of the debugging  
8       information entry.
- 9       4. For each of the attributes in [Table 7.32 on the following page](#) that are present  
10      in the debugging information entry, in the order listed, append to S a marker  
11      letter (see below), the DWARF attribute code, and the attribute value.

12      Note that except for the initial [DW\\_AT\\_name](#) attribute, attributes are  
13      appended in order according to the alphabetical spelling of their identifier.

14      If an implementation defines any vendor-specific attributes, any such  
15      attributes that are essential to the definition of the type are also included at  
16      the end of the above list, in their own alphabetical suborder.

17      An attribute that refers to another type entry T is processed as follows:

- 18      a) If T is in the list V at some V[x], use the letter 'R' as the marker and use  
19      the unsigned LEB128 encoding of x as the attribute value.
- 20      b) Otherwise, append type T to the list V, then use the letter 'T' as the  
21      marker, process the type T recursively by performing Steps 2 through 7,  
22      and use the result as the attribute value.

23      Other attribute values use the letter 'A' as the marker, and the value consists  
24      of the form code (encoded as an unsigned LEB128 value) followed by the  
25      encoding of the value according to the form code. To ensure reproducibility  
26      of the signature, the set of forms used in the signature computation is limited  
27      to the following: [DW\\_FORM\\_sdata](#), [DW\\_FORM\\_flag](#), [DW\\_FORM\\_string](#),  
28      [DW\\_FORM\\_exprval](#), and [DW\\_FORM\\_block](#).



Table 7.32: Attributes used in type signature computation

---

DW_AT_name	DW_AT_endianity
DW_AT_accessibility	DW_AT_enum_class
DW_AT_address_class	DW_AT_explicit
DW_AT_alignment	DW_AT_is_optional
DW_AT_allocated	DW_AT_location
DW_AT_artificial	DW_AT_lower_bound
DW_AT_associated	DW_AT_mutable
DW_AT_binary_scale	DW_AT_ordering
DW_AT_bit_size	DW_AT_picture_string
DW_AT_bit_stride	DW_AT_prototyped
DW_AT_byte_size	DW_AT_rank
DW_AT_byte_stride	DW_AT_reference
DW_AT_const_expr	DW_AT_rvalue_reference
DW_AT_const_value	DW_AT_scale_divisor
DW_AT_containing_type	DW_AT_scale_multiplier
DW_AT_count	DW_AT_small
DW_AT_data_bit_offset	DW_AT_string_length
DW_AT_data_location	DW_AT_string_length_bit_size
DW_AT_data_member_location	DW_AT_string_length_byte_size
DW_AT_decimal_scale	DW_AT_threads_scaled
DW_AT_decimal_sign	DW_AT_upper_bound
DW_AT_default_value	DW_AT_use_location
DW_AT_digit_count	DW_AT_use_UTF8
DW_AT_discr	DW_AT_variable_parameter
DW_AT_discr_list	DW_AT_virtuality
DW_AT_discr_value	DW_AT_visibility
DW_AT_encoding	DW_AT_vtable_elem_location

---

- 1 5. If the tag in Step 3 is one of `DW_TAG_pointer_type`,  
2 `DW_TAG_reference_type`, `DW_TAG_rvalue_reference_type`,  
3 `DW_TAG_ptr_to_member_type`, or `DW_TAG_friend`, and the referenced  
4 type (via the `DW_AT_type` or `DW_AT_friend` attribute) has a `DW_AT_name`  
5 attribute, append to S the letter 'N', the DWARF attribute code (`DW_AT_type`  
6 or `DW_AT_friend`), the context of the type (according to the method in Step  
7 2), the letter 'E', and the name of the type. For `DW_TAG_friend`, if the  
8 referenced entry is a `DW_TAG_subprogram`, the context is omitted and the

## Chapter 7. Data Representation

1 name to be used is the ABI-specific name of the subprogram (for example, the  
2 mangled linker name).

- 3 6. If the tag in Step 3 is not one of [DW\\_TAG\\_pointer\\_type](#),  
4 [DW\\_TAG\\_reference\\_type](#), [DW\\_TAG\\_rvalue\\_reference\\_type](#),  
5 [DW\\_TAG\\_ptr\\_to\\_member\\_type](#), or [DW\\_TAG\\_friend](#), but has a [DW\\_AT\\_type](#)  
6 attribute, or if the referenced type (via the [DW\\_AT\\_type](#) or [DW\\_AT\\_friend](#)  
7 attribute) does not have a [DW\\_AT\\_name](#) attribute, the attribute is processed  
8 according to the method in Step 4 for an attribute that refers to another type  
9 entry.
- 10 7. Visit each child C of the debugging information entry as follows: If C is a  
11 nested type entry or a member function entry, and has a [DW\\_AT\\_name](#)  
12 attribute, append to S the letter 'S', the tag of C, and its name; otherwise,  
13 process C recursively by performing Steps 3 through 7, appending the result  
14 to S. Following the last child (or if there are no children), append a zero byte.

15 For the purposes of this algorithm, if a debugging information entry S has a  
16 [DW\\_AT\\_specification](#) attribute that refers to another entry D (which has a  
17 [DW\\_AT\\_declaration](#) attribute), then S inherits the attributes and children of D,  
18 and S is processed as if those attributes and children were present in the entry S.  
19 Exception: if a particular attribute is found in both S and D, the attribute in S is  
20 used and the corresponding one in D is ignored.

21 DWARF tag and attribute codes are appended to the sequence as unsigned  
22 LEB128 values, using the values defined earlier in this chapter.

23 *A grammar describing this computation may be found in Appendix E.2.2 on page 406.*

24 *An attribute that refers to another type entry is recursively processed or replaced with the*  
25 *name of the referent (in Step 4, 5 or 6). If neither treatment applies to an attribute that*  
26 *references another type entry, the entry that contains that attribute is not suitable for a*  
27 *separate type unit.*

28 *If a debugging information entry contains an attribute from the list above that would*  
29 *require an unsupported form, that entry is not suitable for a separate type unit.*

30 *A type is suitable for a separate type unit only if all of the type entries that it contains or*  
31 *refers to in Steps 6 and 7 are themselves suitable for a separate type unit.*

## Chapter 7. Data Representation

1 *Where the DWARF producer may reasonably choose two or more different forms for a*  
2 *given attribute, it should choose the simplest possible form in computing the signature.*  
3 *(For example, a constant value should be preferred to an expression when possible.)*

4 Once the string S has been formed from the DWARF encoding, an 16-byte MD5  
5 digest is computed for the string and the last eight bytes are taken as the type  
6 signature.

7 *The string S is intended to be a flattened representation of the type that uniquely*  
8 *identifies that type (that is, a different type is highly unlikely to produce the same string).*

9 *A debugging information entry is not placed in a separate type unit if any of the*  
10 *following apply:*

- 11 • *The entry has an attribute whose value is a location description, and the location*  
12 *description contains a reference to another debugging information entry (for*  
13 *example, a `DW_OP_call_ref` operator), as it is unlikely that the entry will remain*  
14 *identical across compilation units.*
- 15 • *The entry has an attribute whose value refers to a code location or a location list.*
- 16 • *The entry has an attribute whose value refers to another debugging information*  
17 *entry that does not represent a type.*

18 *Certain attributes are not included in the type signature:*

- 19 • *The `DW_AT_declaration` attribute is not included because it indicates that the*  
20 *debugging information entry represents an incomplete declaration, and incomplete*  
21 *declarations should not be placed in separate type units.*
- 22 • *The `DW_AT_description` attribute is not included because it does not provide any*  
23 *information unique to the defining declaration of the type.*
- 24 • *The `DW_AT_decl_file`, `DW_AT_decl_line`, and `DW_AT_decl_column` attributes*  
25 *are not included because they may vary from one source file to the next, and would*  
26 *prevent two otherwise identical type declarations from producing the same MD5*  
27 *digest.*
- 28 • *The `DW_AT_object_pointer` attribute is not included because the information it*  
29 *provides is not necessary for the computation of a unique type signature.*

30 *Nested types and some types referred to by a debugging information entry are encoded by*  
31 *name rather than by recursively encoding the type to allow for cases where a complete*  
32 *definition of the type might not be available in all compilation units.*

1 *If a type definition contains the definition of a member function, it cannot be moved as is*  
 2 *into a type unit, because the member function contains attributes that are unique to that*  
 3 *compilation unit. Such a type definition can be moved to a type unit by rewriting the*  
 4 *debugging information entry tree, moving the member function declaration into a*  
 5 *separate declaration tree, and replacing the function definition in the type with a*  
 6 *non-defining declaration of the function (as if the function had been defined out of line).*

7 An example that illustrates the computation of an MD5 digest may be found in  
 8 Appendix E.2 on page 396.

### 9 7.33 Name Table Hash Function

10 The hash function used for hashing name strings in the accelerated access name  
 11 index table (see Section 6.1 on page 141) is defined in C as shown in Figure 7.1  
 12 following.<sup>17</sup>

```

uint32_t /* must be a 32-bit integer type */
hash(unsigned char *str)
{
    uint32_t hash = 5381;
    int c;

    while (c = *str++)
        hash = hash * 33 + c;

    return hash;
}

```

Figure 7.1: Name Table Hash Function Definition

### 13 7.34 Contiguous Tables

14 Tables within each section must be contiguous with the preceding table in that  
 15 section, or the beginning of the section if there is no preceding table.

---

<sup>17</sup> This hash function is sometimes known as the "Bernstein hash function" or the "DJB hash function" (see, for example, [http://en.wikipedia.org/wiki/List\\_of\\_hash\\_functions](http://en.wikipedia.org/wiki/List_of_hash_functions) or <http://stackoverflow.com/questions/10696223/reason-for-5381-number-in-djb-hash-function>).

# Appendix A

## Attributes by Tag Value (Informative)

The table below enumerates the attributes that are most applicable to each type of debugging information entry. DWARF does not in general require that a given debugging information entry contain a particular attribute or set of attributes. Instead, a DWARF producer is free to generate any, all, or none of the attributes described in the text as being applicable to a given entry. Other attributes (both those defined within this document but not explicitly associated with the entry in question, and new, vendor-defined ones) may also appear in a given debugging information entry. Therefore, the table may be taken as instructive, but cannot be considered definitive.

In the following table, the following special conventions apply:

1. The DECL pseudo-attribute stands for all three of the declaration coordinates [DW\\_AT\\_decl\\_column](#), [DW\\_AT\\_decl\\_file](#) and [DW\\_AT\\_decl\\_line](#).
2. The [DW\\_AT\\_description](#) attribute can be used on any debugging information entry that may have a [DW\\_AT\\_name](#) attribute. For simplicity, this attribute is not explicitly shown.
3. The [DW\\_AT\\_sibling](#) attribute can be used on any debugging information entry. For simplicity, this attribute is not explicitly shown.
4. The [DW\\_AT\\_abstract\\_origin](#) attribute can be used with almost any debugging information entry; the exceptions are mostly the compilation unit-like entries. For simplicity, this attribute is not explicitly shown.
5. The [DW\\_AT\\_artificial](#) attribute can be used with any declarative debugging information entry. For simplicity, this attribute is not shown.

## Appendix A. Attributes by Tag (Informative)

Table A.1: Attributes by tag value

TAG name	Applicable attributes
DW_TAG_access_declaration	DECL DW_AT_accessibility DW_AT_name
DW_TAG_array_type	DECL DW_AT_accessibility DW_AT_alignment DW_AT_allocated DW_AT_associated DW_AT_bit_size DW_AT_bit_stride DW_AT_byte_size DW_AT_data_location DW_AT_declaration DW_AT_name DW_AT_ordering DW_AT_rank DW_AT_specification DW_AT_start_scope DW_AT_type DW_AT_visibility
DW_TAG_atomic_type	DECL DW_AT_alignment DW_AT_name DW_AT_type

*Continued on next page*

## Appendix A. Attributes by Tag (Informative)

TAG name	Applicable attributes
DW_TAG_base_type	DECL DW_AT_alignment DW_AT_allocated DW_AT_associated DW_AT_bias DW_AT_binary_scale DW_AT_bit_size DW_AT_byte_size DW_AT_data_bit_offset DW_AT_data_location DW_AT_decimal_scale DW_AT_decimal_sign DW_AT_digit_count DW_AT_encoding DW_AT_endianity DW_AT_name DW_AT_picture_string DW_AT_scale_divisor DW_AT_scale_multiplier DW_AT_small
DW_TAG_call_site	DW_AT_call_column DW_AT_call_file DW_AT_call_line DW_AT_call_origin DW_AT_call_pc DW_AT_call_return_pc DW_AT_call_tail_call DW_AT_call_target DW_AT_call_target_clobbered DW_AT_type

*Continued on next page*

Appendix A. Attributes by Tag (Informative)

<b>TAG name</b>	<b>Applicable attributes</b>
DW_TAG_call_site_parameter	DW_AT_call_data_location DW_AT_call_data_value DW_AT_call_parameter DW_AT_call_value DW_AT_location DW_AT_name DW_AT_type
DW_TAG_catch_block	DECL DW_AT_entry_pc DW_AT_high_pc DW_AT_low_pc DW_AT_ranges
DW_TAG_class_type	DECL DW_AT_accessibility DW_AT_alignment DW_AT_allocated DW_AT_associated DW_AT_bit_size DW_AT_byte_size DW_AT_calling_convention DW_AT_data_location DW_AT_declaration DW_AT_export_symbols DW_AT_name DW_AT_signature DW_AT_specification DW_AT_start_scope DW_AT_visibility
DW_TAG_coarray_type	DECL DW_AT_alignment DW_AT_bit_size DW_AT_byte_size DW_AT_name DW_AT_type

*Continued on next page*



## Appendix A. Attributes by Tag (Informative)

TAG name	Applicable attributes
DW_TAG_common_block	DECL DW_AT_declaration DW_AT_linkage_name DW_AT_location DW_AT_name DW_AT_visibility
DW_TAG_common_inclusion	DECL DW_AT_common_reference DW_AT_declaration DW_AT_visibility
DW_TAG_compile_unit	DW_AT_addr_base DW_AT_base_types DW_AT_comp_dir DW_AT_entry_pc DW_AT_identifier_case DW_AT_high_pc DW_AT_language_name DW_AT_language_version DW_AT_low_pc DW_AT_macros DW_AT_main_subprogram DW_AT_name DW_AT_producer DW_AT_ranges DW_AT_rnglists_base DW_AT_stmt_list DW_AT_str_offsets DW_AT_use_UTF8
DW_TAG_condition	DECL DW_AT_name
DW_TAG_const_type	DECL DW_AT_alignment DW_AT_name DW_AT_type

*Continued on next page*

## Appendix A. Attributes by Tag (Informative)

TAG name	Applicable attributes
DW_TAG_constant	DECL DW_AT_accessibility DW_AT_const_value DW_AT_declaration DW_AT_endianity DW_AT_external DW_AT_linkage_name DW_AT_name DW_AT_start_scope DW_AT_type DW_AT_visibility
DW_TAG_dwarf_procedure	DW_AT_location
DW_TAG_dynamic_type	DECL DW_AT_alignment DW_AT_allocated DW_AT_associated DW_AT_data_location DW_AT_name DW_AT_type
DW_TAG_entry_point	DECL DW_AT_address_class DW_AT_frame_base DW_AT_linkage_name DW_AT_low_pc DW_AT_name DW_AT_return_addr DW_AT_static_link DW_AT_type

*Continued on next page*

## Appendix A. Attributes by Tag (Informative)

TAG name	Applicable attributes
DW_TAG_enumeration_type	DECL DW_AT_accessibility DW_AT_alignment DW_AT_allocated DW_AT_associated DW_AT_bit_size DW_AT_bit_stride DW_AT_byte_size DW_AT_byte_stride DW_AT_data_location DW_AT_declaration DW_AT_enum_class DW_AT_name DW_AT_signature DW_AT_specification DW_AT_start_scope DW_AT_type DW_AT_visibility
DW_TAG_enumerator	DECL DW_AT_const_value DW_AT_name
DW_TAG_file_type	DECL DW_AT_alignment DW_AT_allocated DW_AT_associated DW_AT_bit_size DW_AT_byte_size DW_AT_data_location DW_AT_name DW_AT_start_scope DW_AT_type DW_AT_visibility

*Continued on next page*

## Appendix A. Attributes by Tag (Informative)

TAG name	Applicable attributes
DW_TAG_formal_parameter	DECL DW_AT_const_value DW_AT_default_value DW_AT_endianity DW_AT_is_optional DW_AT_location DW_AT_name DW_AT_type DW_AT_variable_parameter
DW_TAG_friend	DECL DW_AT_friend
DW_TAG_generic_subrange	DECL DW_AT_accessibility DW_AT_alignment DW_AT_allocated DW_AT_associated DW_AT_bit_size DW_AT_bit_stride DW_AT_byte_size DW_AT_byte_stride DW_AT_count DW_AT_data_location DW_AT_declaration DW_AT_lower_bound DW_AT_name DW_AT_threads_scaled DW_AT_type DW_AT_upper_bound DW_AT_visibility
DW_TAG_immutable_type	DECL DW_AT_name DW_AT_type

*Continued on next page*

## Appendix A. Attributes by Tag (Informative)

TAG name	Applicable attributes
DW_TAG_imported_declaration	DECL DW_AT_accessibility DW_AT_import DW_AT_name DW_AT_start_scope
DW_TAG_imported_module	DECL DW_AT_import DW_AT_start_scope
DW_TAG_imported_unit	DW_AT_import
DW_TAG_inheritance	DECL DW_AT_accessibility DW_AT_data_member_location DW_AT_type DW_AT_virtuality
DW_TAG_inlined_subroutine	DW_AT_call_column DW_AT_call_file DW_AT_call_line DW_AT_const_expr DW_AT_entry_pc DW_AT_high_pc DW_AT_low_pc DW_AT_ranges DW_AT_return_addr DW_AT_start_scope DW_AT_trampoline
DW_TAG_interface_type	DECL DW_AT_accessibility DW_AT_alignment DW_AT_name DW_AT_signature DW_AT_start_scope

*Continued on next page*

Appendix A. Attributes by Tag (Informative)

<b>TAG name</b>	<b>Applicable attributes</b>
DW_TAG_label	DECL DW_AT_low_pc DW_AT_name DW_AT_start_scope
DW_TAG_lexical_block	DECL DW_AT_entry_pc DW_AT_high_pc DW_AT_low_pc DW_AT_name DW_AT_ranges
DW_TAG_member	DECL DW_AT_accessibility DW_AT_bit_size DW_AT_byte_size DW_AT_data_bit_offset DW_AT_data_member_location DW_AT_declaration DW_AT_mutable DW_AT_name DW_AT_type DW_AT_visibility
DW_TAG_module	DECL DW_AT_accessibility DW_AT_declaration DW_AT_entry_pc DW_AT_high_pc DW_AT_low_pc DW_AT_name DW_AT_priority DW_AT_ranges DW_AT_specification DW_AT_visibility

*Continued on next page*

## Appendix A. Attributes by Tag (Informative)

TAG name	Applicable attributes
DW_TAG_namelist	DECL DW_AT_accessibility DW_AT_declaration DW_AT_name DW_AT_visibility
DW_TAG_namelist_item	DECL DW_AT_namelist_item
DW_TAG_namespace	DECL DW_AT_export_symbols DW_AT_extension DW_AT_name DW_AT_start_scope
DW_TAG_packed_type	DECL DW_AT_alignment DW_AT_name DW_AT_type

*Continued on next page*

## Appendix A. Attributes by Tag (Informative)

TAG name	Applicable attributes
DW_TAG_partial_unit	DW_AT_addr_base DW_AT_base_types DW_AT_comp_dir DW_AT_dwo_name DW_AT_entry_pc DW_AT_identifier_case DW_AT_high_pc DW_AT_language_name DW_AT_language_version DW_AT_low_pc DW_AT_macros DW_AT_main_subprogram DW_AT_name DW_AT_producer DW_AT_ranges DW_AT_rnglists_base DW_AT_stmt_list DW_AT_str_offsets DW_AT_use_UTF8
DW_TAG_pointer_type	DECL DW_AT_address_class DW_AT_alignment DW_AT_bit_size DW_AT_byte_size DW_AT_name DW_AT_type

*Continued on next page*



Appendix A. Attributes by Tag (Informative)

<b>TAG name</b>	<b>Applicable attributes</b>
DW_TAG_ptr_to_member_type	DECL DW_AT_address_class DW_AT_alignment DW_AT_allocated DW_AT_associated DW_AT_containing_type DW_AT_data_location DW_AT_declaration DW_AT_name DW_AT_type DW_AT_use_location DW_AT_visibility
DW_TAG_reference_type	DECL DW_AT_address_class DW_AT_alignment DW_AT_bit_size DW_AT_byte_size DW_AT_name DW_AT_type
DW_TAG_restrict_type	DECL DW_AT_alignment DW_AT_name DW_AT_type
DW_TAG_rvalue_reference_type	DECL DW_AT_address_class DW_AT_alignment DW_AT_bit_size DW_AT_byte_size DW_AT_name DW_AT_type

*Continued on next page*

## Appendix A. Attributes by Tag (Informative)

TAG name	Applicable attributes
DW_TAG_set_type	DECL DW_AT_accessibility DW_AT_alignment DW_AT_allocated DW_AT_associated DW_AT_bit_size DW_AT_byte_size DW_AT_data_location DW_AT_declaration DW_AT_name DW_AT_start_scope DW_AT_type DW_AT_visibility
DW_TAG_shared_type	DECL DW_AT_count DW_AT_alignment DW_AT_name DW_AT_type
DW_TAG_skeleton_unit	DW_AT_addr_base DW_AT_comp_dir DW_AT_dwo_name DW_AT_high_pc DW_AT_low_pc DW_AT_ranges DW_AT_rnglists_base DW_AT_stmt_list DW_AT_str_offsets DW_AT_use_UTF8

*Continued on next page*

## Appendix A. Attributes by Tag (Informative)

TAG name	Applicable attributes
DW_TAG_string_type	DECL DW_AT_alignment DW_AT_accessibility DW_AT_allocated DW_AT_associated DW_AT_bit_size DW_AT_byte_size DW_AT_data_location DW_AT_declaration DW_AT_name DW_AT_start_scope DW_AT_string_length DW_AT_string_length_bit_size DW_AT_string_length_byte_size DW_AT_visibility
DW_TAG_structure_type	DECL DW_AT_accessibility DW_AT_alignment DW_AT_allocated DW_AT_associated DW_AT_bit_size DW_AT_byte_size DW_AT_calling_convention DW_AT_data_location DW_AT_declaration DW_AT_export_symbols DW_AT_name DW_AT_signature DW_AT_specification DW_AT_start_scope DW_AT_visibility

*Continued on next page*

## Appendix A. Attributes by Tag (Informative)

TAG name	Applicable attributes
DW_TAG_subprogram	DECL DW_AT_accessibility DW_AT_address_class DW_AT_alignment DW_AT_calling_convention DW_AT_declaration DW_AT_defaulted DW_AT_deleted DW_AT_elemental DW_AT_entry_pc DW_AT_explicit DW_AT_external DW_AT_frame_base DW_AT_high_pc DW_AT_inline DW_AT_linkage_name DW_AT_low_pc DW_AT_main_subprogram DW_AT_name DW_AT_noreturn DW_AT_object_pointer DW_AT_prototyped DW_AT_pure DW_AT_ranges DW_AT_recursive DW_AT_reference DW_AT_return_addr DW_AT_rvalue_reference DW_AT_specification <i>Additional attributes continue on next page</i>

*Continued on next page*

## Appendix A. Attributes by Tag (Informative)

TAG name	Applicable attributes
DW_TAG_subprogram (cont.)	DW_AT_start_scope DW_AT_static_link DW_AT_trampoline DW_AT_type DW_AT_visibility DW_AT_virtuality DW_AT_vtable_elem_location
DW_TAG_subrange_type	DECL DW_AT_accessibility DW_AT_alignment DW_AT_allocated DW_AT_associated DW_AT_bit_size DW_AT_bit_stride DW_AT_byte_size DW_AT_byte_stride DW_AT_count DW_AT_data_location DW_AT_declaration DW_AT_lower_bound DW_AT_name DW_AT_threads_scaled DW_AT_type DW_AT_upper_bound DW_AT_visibility

*Continued on next page*

Appendix A. Attributes by Tag (Informative)

<b>TAG name</b>	<b>Applicable attributes</b>
DW_TAG_subroutine_type	DECL DW_AT_accessibility DW_AT_address_class DW_AT_alignment DW_AT_allocated DW_AT_associated DW_AT_data_location DW_AT_declaration DW_AT_name DW_AT_prototyped DW_AT_reference DW_AT_rvalue_reference DW_AT_start_scope DW_AT_type DW_AT_visibility
DW_TAG_template_alias	DECL DW_AT_accessibility DW_AT_allocated DW_AT_associated DW_AT_data_location DW_AT_declaration DW_AT_name DW_AT_signature DW_AT_start_scope DW_AT_type DW_AT_visibility
DW_TAG_template_type_parameter	DECL DW_AT_default_value DW_AT_name DW_AT_type

*Continued on next page*

Appendix A. Attributes by Tag (Informative)

<b>TAG name</b>	<b>Applicable attributes</b>
DW_TAG_template_value_parameter	DECL DW_AT_const_value DW_AT_default_value DW_AT_name DW_AT_type
DW_TAG_thrown_type	DECL DW_AT_alignment DW_AT_allocated DW_AT_associated DW_AT_data_location DW_AT_name DW_AT_type
DW_TAG_try_block	DECL DW_AT_entry_pc DW_AT_high_pc DW_AT_low_pc DW_AT_ranges
DW_TAG_typedef	DECL DW_AT_accessibility DW_AT_alignment DW_AT_allocated DW_AT_associated DW_AT_data_location DW_AT_declaration DW_AT_name DW_AT_start_scope DW_AT_type DW_AT_visibility
DW_TAG_type_unit	DW_AT_language_name DW_AT_language_version DW_AT_stmt_list DW_AT_str_offsets DW_AT_use_UTF8

*Continued on next page*

## Appendix A. Attributes by Tag (Informative)

TAG name	Applicable attributes
DW_TAG_union_type	DECL DW_AT_accessibility DW_AT_alignment DW_AT_allocated DW_AT_associated DW_AT_bit_size DW_AT_byte_size DW_AT_calling_convention DW_AT_data_location DW_AT_declaration DW_AT_export_symbols DW_AT_name DW_AT_signature DW_AT_specification DW_AT_start_scope DW_AT_visibility
DW_TAG_unspecified_parameters	DECL
DW_TAG_unspecified_type	DECL DW_AT_name
DW_TAG_variable	DECL DW_AT_accessibility DW_AT_alignment DW_AT_const_expr DW_AT_const_value DW_AT_declaration DW_AT_endianity DW_AT_external DW_AT_linkage_name DW_AT_location DW_AT_name DW_AT_specification DW_AT_start_scope DW_AT_type DW_AT_visibility

*Continued on next page*



Appendix A. Attributes by Tag (Informative)

<b>TAG name</b>	<b>Applicable attributes</b>
DW_TAG_variant	DECL DW_AT_accessibility DW_AT_declaration DW_AT_discr_list DW_AT_discr_value
DW_TAG_variant_part	DECL DW_AT_accessibility DW_AT_declaration DW_AT_discr DW_AT_type
DW_TAG_volatile_type	DECL DW_AT_name DW_AT_type
DW_TAG_with_stmt	DECL DW_AT_accessibility DW_AT_address_class DW_AT_declaration DW_AT_entry_pc DW_AT_high_pc DW_AT_location DW_AT_low_pc DW_AT_ranges DW_AT_type DW_AT_visibility

# Appendix B

## Debug Section Relationships (Informative)

DWARF information is organized into multiple program sections, each of which holds a particular kind of information. In some cases, information in one section refers to information in one or more of the others. These relationships are illustrated by the diagrams and associated notes on the following pages.

In the figures, a section is shown as a shaded oval with the name of the section inside. References from one section to another are shown by an arrow. In the first figure, the arrow is annotated with an unshaded box which contains an indication of the construct (such as an attribute or form) that encodes the reference. In the second figure, this box is left out for reasons of space in favor of a label annotation that is explained in the subsequent notes.

### B.1 Normal DWARF Section Relationships

Figure B.1 following illustrates the DWARF section relations without split DWARF object files involved. Similarly, it does not show the relationships between the main debugging sections of an executable or sharable file and a related supplementary object file.

## 1 **B.2 Split DWARF Section Relationships**

2 Figure [B.2 on page 290](#) illustrates the DWARF section relationships for split  
3 DWARF object files. However, it does not show the relationships between the  
4 main debugging sections of an executable or shareable file and a related  
5 supplementary object file. For space reasons, the figure omits some details that  
6 are shown in Figure [B.1](#), such as indirect references using indexing sections (such  
7 as `.debug_str_offsets`).

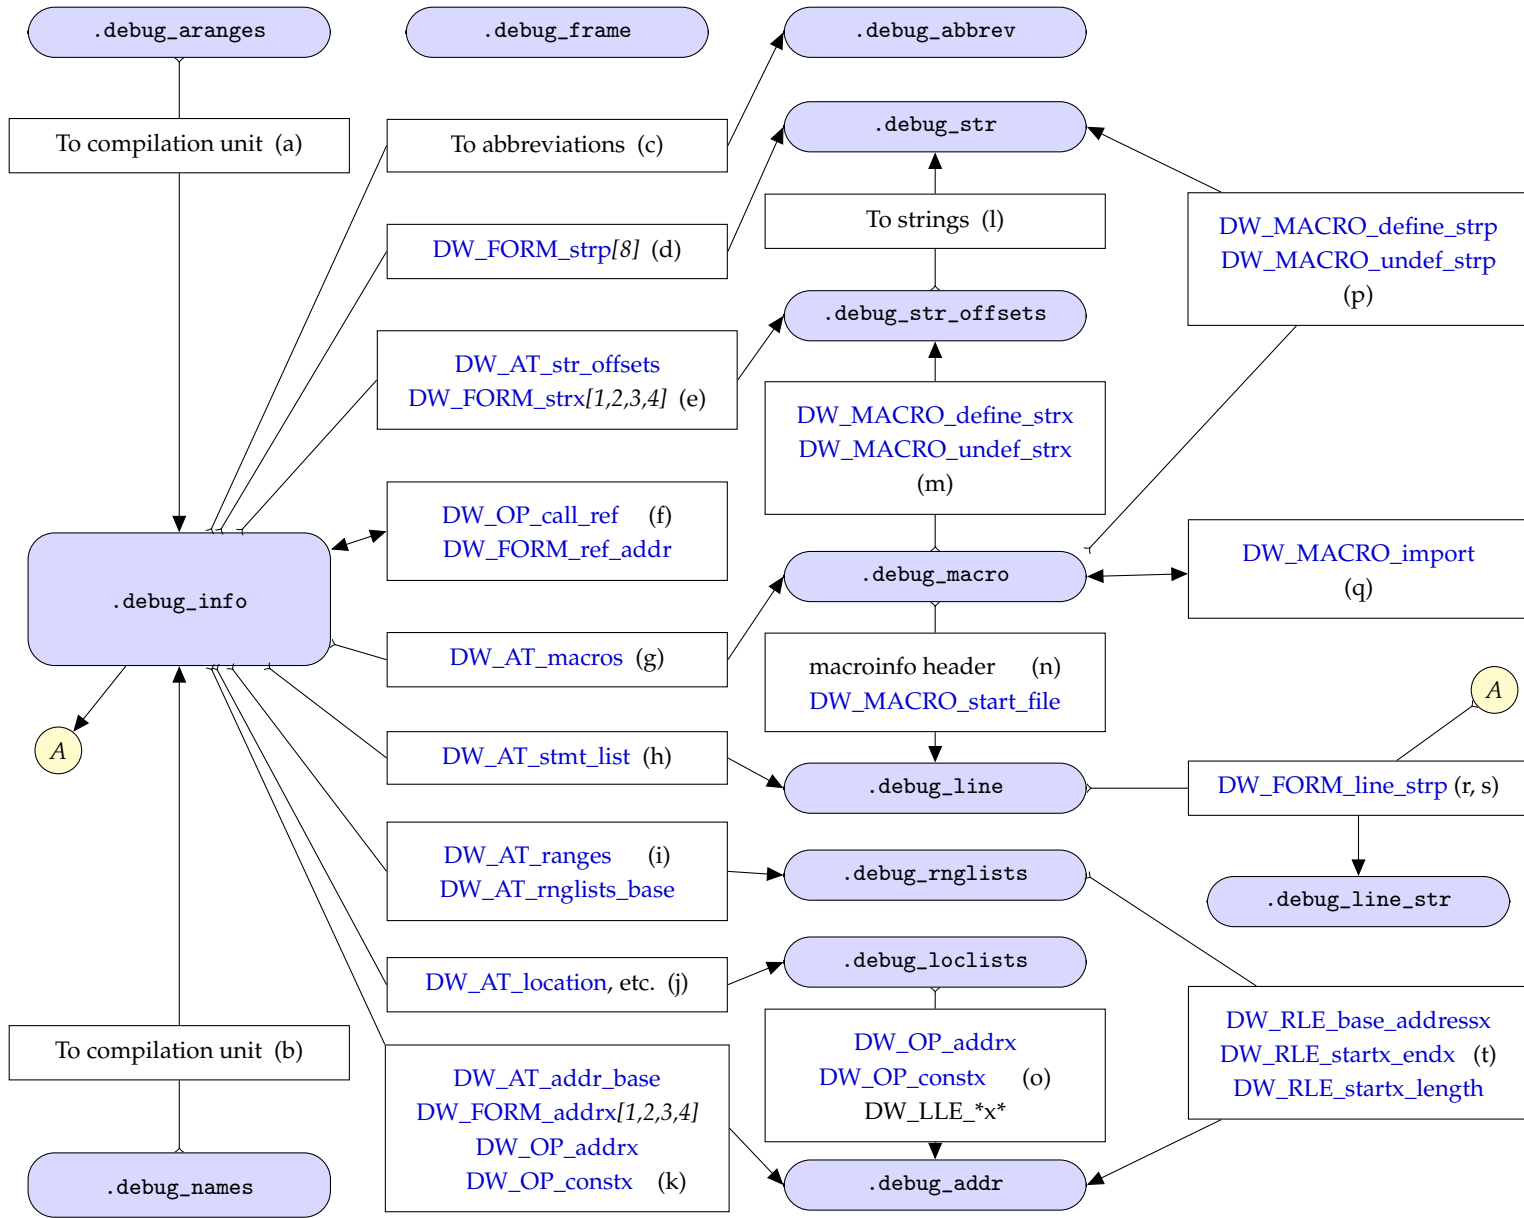


Figure B.1: Debug section relationships

## Appendix B. Debug Section Relationships (Informative)

### Notes for Figure B.1

- 1
- 2     **(a)** `.debug_aranges` to `.debug_info`  
3         The `debug_info_offset` value in the header is the offset in the `.debug_info`  
4         section of the corresponding compilation unit header (not the compilation  
5         unit entry).
- 6     **(b)** `.debug_names` to `.debug_info`  
7         The list of compilation units following the header contains the offsets in the  
8         `.debug_info` section of the corresponding compilation unit headers (not the  
9         compilation unit entries).
- 10    **(c)** `.debug_info` to `.debug_abbrev`  
11         The `debug_abbrev_offset` value in the header is the offset in the  
12         `.debug_abbrev` section of the abbreviations for that compilation unit.
- 13    **(d)** `.debug_info` to `.debug_str`  
14         Attribute values of class string may have form `DW_FORM_strp` or  
15         `DW_FORM_strp8`, whose value is the offset in the `.debug_str` section of  
16         the corresponding string.
- 17    **(e)** `.debug_info` to `.debug_str_offsets`  
18         The value of the `DW_AT_str_offsets` attribute in a `DW_TAG_compile_unit`,  
19         `DW_TAG_type_unit` or `DW_TAG_partial_unit` DIE is the offset in the  
20         `.debug_str_offsets` section of the header of the string offsets information  
21         for that unit. In addition, attribute values of class string may have one of  
22         the forms `DW_FORM_strx`, `DW_FORM_strx1`, `DW_FORM_strx2`,  
23         `DW_FORM_strx3` or `DW_FORM_strx4`, whose value is an index into the  
24         string offsets table.
- 25    **(f)** `.debug_info` to `.debug_info`  
26         The operand of the `DW_OP_call_ref` DWARF expression operator is the  
27         offset of a debugging information entry in the `.debug_info` section of  
28         another compilation. Similarly for attribute operands that use  
29         `DW_FORM_ref_addr`.
- 30    **(g)** `.debug_info` to `.debug_macro`  
31         An attribute value of class `macptr` (specifically form `DW_FORM_sec_offset`)  
32         is an offset within the `.debug_macro` section of the beginning of the macro  
33         information for the referencing unit.
- 34    **(h)** `.debug_info` to `.debug_line`  
35         An attribute value of class `lineptr` (specifically form `DW_FORM_sec_offset`)  
36         is an offset in the `.debug_line` section of the beginning of the line number  
37         information for the referencing unit.

## Appendix B. Debug Section Relationships (Informative)

- 1       **(i)** `.debug_info` to `.debug_rnglists`  
2       An attribute value of class `rnglist` (specifically form `DW_FORM_rnglistx` or  
3       `DW_FORM_sec_offset`) is an index or offset within the `.debug_rnglists`  
4       section of a range list.
- 5       **(j)** `.debug_info` to `.debug_loclists`  
6       An attribute value of class `loclist` (specifically form `DW_FORM_loclistx` or  
7       `DW_FORM_sec_offset`) is an index or offset within the `.debug_loclists`  
8       section of a location list.
- 9       **(k)** `.debug_info` to `.debug_addr`  
10       The value of the `DW_AT_addr_base` attribute in the  
11       `DW_TAG_compile_unit` or `DW_TAG_partial_unit` DIE is the offset in the  
12       `.debug_addr` section of the machine addresses for that unit.  
13       `DW_FORM_addrx`, `DW_FORM_addrx1`, `DW_FORM_addrx2`,  
14       `DW_FORM_addrx3`, `DW_FORM_addrx4`, `DW_OP_addrx` and  
15       `DW_OP_constx` contain indices relative to that offset.
- 16       **(l)** `.debug_str_offsets` to `.debug_str`  
17       Entries in the string offsets table are offsets to the corresponding string text  
18       in the `.debug_str` section.
- 19       **(m)** `.debug_macro` to `.debug_str_offsets`  
20       The second operand of a `DW_MACRO_define_strx` or  
21       `DW_MACRO_undef_strx` macro information entry is an index into the  
22       string offset table in the `.debug_str_offsets` section.
- 23       **(n)** `.debug_macro` to `.debug_line`  
24       The second operand of `DW_MACRO_start_file` refers to a file entry in the  
25       `.debug_line` section relative to the start of that section given in the macro  
26       information header.
- 27       **(o)** `.debug_loclists` to `.debug_addr`  
28       `DW_OP_addrx` and `DW_OP_constx` operators that occur in the  
29       `.debug_loclists` section refer indirectly to the `.debug_addr` section by way  
30       of the `DW_AT_addr_base` attribute in the associated `.debug_info` section.  
31       Also, some operands of the `DW_LLE_base_addressx`,  
32       `DW_LLE_startx_endx` and `DW_LLE_startx_length` location list entries have  
33       operands that are an index into the `.debug_addr` section.
- 34       **(p)** `.debug_macro` to `.debug_str`  
35       The second operand of a `DW_MACRO_define_strp` or  
36       `DW_MACRO_undef_strp` macro information entry is an index into the  
37       string table in the `.debug_str` section.

## Appendix B. Debug Section Relationships (Informative)

- 1       **(q)** `.debug_macro` **to** `.debug_macro`  
2           The operand of a `DW_MACRO_import` macro information entry is an  
3           offset into another part of the `.debug_macro` section to the header for the  
4           sequence to be replicated.
- 5       **(r)** `.debug_line` **to** `.debug_line_str`  
6           The value of a `DW_FORM_line_strp` form refers to a string section specific  
7           to the line number table. This form can be used in a `.debug_line` section (as  
8           well as in a `.debug_info` section).
- 9       **(s)** `.debug_info` **to** `.debug_line_str`  
10          The value of a `DW_FORM_line_strp` form refers to a string section specific  
11          to the line number table. This form can be used in a `.debug_info` section (as  
12          well as in a `.debug_line` section).<sup>1</sup>
- 13       **(t)** `.debug_rnglists` **to** `.debug_addr`  
14          Some operands of `DW_RLE_base_addressx`, `DW_RLE_startx_endx` and  
15          `DW_RLE_startx_length` range list entries are an an index into the  
16          `.debug_addr` section.

---

<sup>1</sup>The circled (A) of the left connects to the circled (A) on the right via hyperspace (a wormhole).

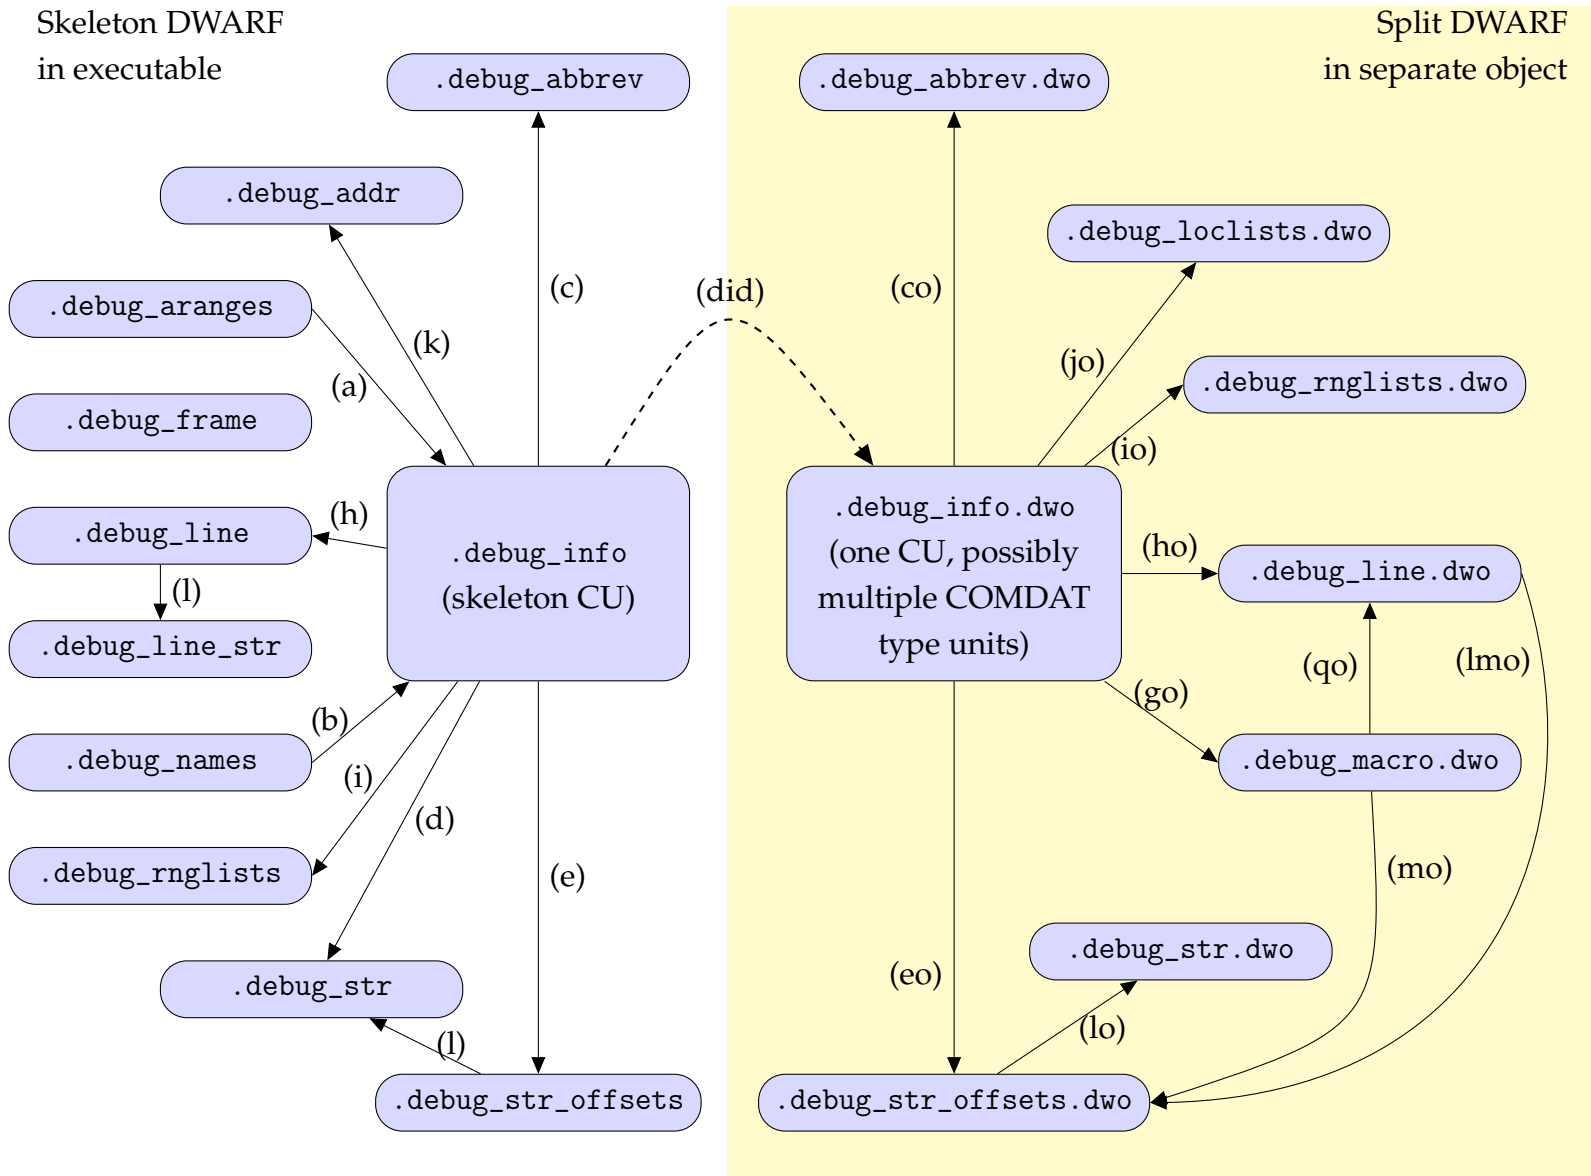


Figure B.2: Split DWARF section relationships



## Appendix B. Debug Section Relationships (Informative)

### Notes for Figure B.2

- 1
- 2     **(a)** `.debug_aranges` to `.debug_info`  
3         The `debug_info_offset` field in the header is the offset in the `.debug_info`  
4         section of the corresponding compilation unit header of the skeleton  
5         `.debug_info` section (not the compilation unit entry). The  
6         [DW\\_AT\\_dwo\\_name](#) attribute in the `.debug_info` skeleton connects the  
7         ranges to the full compilation unit in `.debug_info.dwo`.
- 8     **(b)** `.debug_names` to `.debug_info`  
9         The `.debug_names` section offsets lists provide an offset for the skeleton  
10         compilation unit and eight byte signatures for the type units that appear  
11         only in the `.debug_info.dwo`. The DIE offsets for these compilation units  
12         and type units refer to the DIEs in the `.debug_info.dwo` section for the  
13         respective compilation unit and type units.
- 14     **(c)** `.debug_info skeleton` to `.debug_abbrev`  
15         The `debug_abbrev_offset` value in the header is the offset in the  
16         `.debug_abbrev` section of the abbreviations for that compilation unit  
17         skeleton.
- 18     **(co)** `.debug_info.dwo` to `.debug_abbrev.dwo`  
19         The `debug_abbrev_offset` value in the header is the offset in the  
20         `.debug_abbrev.dwo` section of the abbreviations for that compilation unit.
- 21     **(d)** `.debug_info` to `.debug_str`  
22         Attribute values of class string may have form [DW\\_FORM\\_strp](#) or  
23         [DW\\_FORM\\_strp8](#), whose value is an offset in the `.debug_str` section of the  
24         corresponding string.
- 25     **(did)** `.debug_info` to `.debug_info.dwo`  
26         The [DW\\_AT\\_dwo\\_name](#) attribute in a skeleton unit identifies the file  
27         containing the corresponding `.dwo` (split) data.
- 28     **(e)** `.debug_info` to `.debug_str_offsets`  
29         Attribute values of class string may have one of the forms [DW\\_FORM\\_strx](#),  
30         [DW\\_FORM\\_strx1](#), [DW\\_FORM\\_strx2](#), [DW\\_FORM\\_strx3](#) or  
31         [DW\\_FORM\\_strx4](#), whose value is an index into the `.debug_str_offsets`  
32         section for the corresponding string.
- 33     **(eo)** `.debug_info.dwo` to `.debug_str_offsets.dwo`  
34         Attribute values of class string may have one of the forms [DW\\_FORM\\_strx](#),  
35         [DW\\_FORM\\_strx1](#), [DW\\_FORM\\_strx2](#), [DW\\_FORM\\_strx3](#) or  
36         [DW\\_FORM\\_strx4](#), whose value is an index into the  
37         `.debug_str_offsets.dwo` section for the corresponding string.

## Appendix B. Debug Section Relationships (Informative)

- 1       **(go)** `.debug_info.dwo` to `.debug_macro.dwo`  
2       An attribute of class `macptr` (specifically `DW_AT_macros` with form  
3       `DW_FORM_sec_offset`) is an offset within the `.debug_macro.dwo` section of  
4       the beginning of the macro information for the referencing unit.
- 5       **(h)** `.debug_info (skeleton)` to `.debug_line`  
6       An attribute value of class `lineptr` (specifically `DW_AT_stmt_list` with form  
7       `DW_FORM_sec_offset`) is an offset within the `.debug_line` section of the  
8       beginning of the line number information for the referencing unit.
- 9       **(ho)** `.debug_info.dwo` to `.debug_line.dwo (skeleton)`  
10       An attribute value of class `lineptr` (specifically `DW_AT_stmt_list` with form  
11       `DW_FORM_sec_offset`) is an offset within the `.debug_line.dwo` section of  
12       the beginning of the line number header information for the referencing  
13       unit (the line table details are not in `.debug_line.dwo` but the line header  
14       with its list of file names is present).
- 15       **(i)** `.debug_info` to `.debug_rnglists`  
16       An attribute value of class `rnglist` (specifically form `DW_FORM_rnglistx` or  
17       `DW_FORM_sec_offset`) is an index or offset within the `.debug_rnglists`  
18       section of a range list.
- 19       **(io)** `.debug_info.dwo` to `.debug_rnglists.dwo`  
20       An attribute value of class `rnglist` (specifically `DW_AT_ranges` with form  
21       `DW_FORM_rnglistx` or `DW_FORM_sec_offset`) is an index or offset within  
22       the `.debug_rnglists.dwo` section of a range list. The format of  
23       `.debug_rnglists.dwo` location list entries is restricted to a subset of those  
24       in `.debug_rnglists`. See Section 2.17.3 on page 53 for details.
- 25       **(jo)** `.debug_info.dwo` to `.debug_loclists.dwo`  
26       An attribute value of class `loclist` (specifically with form  
27       `DW_FORM_loclistx` or `DW_FORM_sec_offset`) is an index or offset within  
28       the `.debug_loclists.dwo` section of a location list. The format of  
29       `.debug_loclists.dwo` location list entries is restricted to a subset of those  
30       in `.debug_loclists`. See Section 2.6.2 on page 43 for details.
- 31       **(k)** `.debug_info` to `.debug_addr`  
32       The value of the `DW_AT_addr_base` attribute in the  
33       `DW_TAG_compile_unit`, `DW_TAG_partial_unit` or `DW_TAG_type_unit`  
34       DIE is the offset in the `.debug_addr` section of the machine addresses for  
35       that unit. `DW_FORM_addrx`, `DW_FORM_addrx1`, `DW_FORM_addrx2`,  
36       `DW_FORM_addrx3`, `DW_FORM_addrx4`, `DW_OP_addrx` and  
37       `DW_OP_constx` contain indices relative to that offset.

## Appendix B. Debug Section Relationships (Informative)

- 1       **(lmo)** `.debug_line.dwo` **to** `.debug_str_offsets.dwo`  
2           The value of a `DW_FORM_line_strp` form refers to a string section specific  
3           to the line number table. This form can be used in a `.debug_line.dwo`  
4           section (as well as in a `.debug_info.dwo` section).
- 5       **(lo)** `.debug_str_offsets.dwo` **to** `.debug_str.dwo`  
6           Entries in the string offsets table are offsets to the corresponding string text  
7           in the `.debug_str.dwo` section.
- 8       **(mo)** `.debug_macro.dwo` **to** `.debug_str_offsets.dwo`  
9           Within the `.debug_macro.dwo` sections, the second operand of  
10          `DW_MACRO_define_strx` and `DW_MACRO_undef_strx` operations is an  
11          unsigned LEB128 value interpreted as an index into the  
12          `.debug_str_offsets.dwo` section.
- 13       **(qo)** `.debug_macro.dwo` **to** `.debug_line.dwo`  
14          Within the `.debug_macro.dwo` sections, if a `DW_MACRO_start_file` entry is  
15          present, the macro header contains an offset into the `.debug_line.dwo`  
16          section.

# Appendix C

## Variable Length Data: Encoding/Decoding (Informative)

Here are algorithms expressed in a C-like pseudo-code to encode and decode signed and unsigned numbers in LEB128 representation.

The encode and decode algorithms given here do not take account of C/C++ rules that mean that in  $E1 < E2$  the type of  $E1$  should be a sufficiently large unsigned type to hold the correct mathematical result. The decode algorithms do not take account of or protect from possibly invalid LEB values, such as values that are too large to fit in the target type or that lack a proper terminator byte. Implementation languages may have additional or different rules.

```
do
{
    byte = low order 7 bits of value;
    value >>= 7;
    if (value != 0) /* more bytes to come */
        set high order bit of byte;
    emit byte;
} while (value != 0);
```

Figure C.1: Algorithm to encode an unsigned integer

## Appendix C. Encoding/Decoding (Informative)

```
more = 1;
negative = (value < 0);
size = no. of bits in signed integer;
while(more)
{
    byte = low order 7 bits of value;
    value >>= 7;
    /* the following is unnecessary if the
     * implementation of >>= uses an arithmetic rather
     * than logical shift for a signed left operand
     */
    if (negative)
        /* sign extend */
        value |= - (1 <<(size - 7));
    /* sign bit of byte is second high order bit (0x40) */
    if ((value == 0 && sign bit of byte is clear) ||
        (value == -1 && sign bit of byte is set))
        more = 0;
    else
        set high order bit of byte;
    emit byte;
}
```

Figure C.2: Algorithm to encode a signed integer

```
result = 0;
shift = 0;
while(true)
{
    byte = next byte in input;
    result |= (low order 7 bits of byte << shift);
    if (high order bit of byte == 0)
        break;
    shift += 7;
}
```

Figure C.3: Algorithm to decode an unsigned LEB128 integer

## Appendix C. Encoding/Decoding (Informative)

```
result = 0;
shift = 0;
size = number of bits in signed integer;
while(true)
{
    byte = next byte in input;
    result |= (low order 7 bits of byte << shift);
    shift += 7;
    /* sign bit of byte is second high order bit (0x40) */
    if (high order bit of byte == 0)
        break;
}
if ((shift < size) && (sign bit of byte is set))
    /* sign extend */
    result |= - (1 << shift);
```

Figure C.4: Algorithm to decode a signed LEB128 integer

Appendix C. Encoding/Decoding (Informative)

*(empty page)*

# Appendix D

## Examples (Informative)

The following sections provide examples that illustrate various aspects of the DWARF debugging information format.

### D.1 General Description Examples

#### D.1.1 Compilation Units and Abbreviations Table Example

Figure [D.1 on the next page](#) depicts the relationship of the abbreviations tables contained in the `.debug_abbrev` section to the information contained in the `.debug_info` section. Values are given in symbolic form, where possible.

The figure corresponds to the following two trivial source files:

File myfile.c

```
typedef char* POINTER;
```

File myfile2.c

```
typedef char* strp;
```



## Appendix D. Examples (Informative)

### Compilation Unit #1: .debug\_info

```

length
4
a1 (abbreviations table offset)
4
-----
1
"myfile.c"
"Best Compiler Corp, V1.3"
"/home/mydir/src"
DW_LNAME_C
0x0
0x55
DW_FORM_sec_offset
0x0
-----
e1: 2
"char"
DW_ATE_unsigned_char
1
-----
e2: 3
e1 (debug info offset)
-----
4
"POINTER"
e2 (debug info offset)
-----
0

```

### Compilation Unit #2: .debug\_info

```

length
4
a1 (abbreviations table offset)
4
-----
...
-----
4
"strp"
e2 (debug info offset)
-----
...

```

### Abbreviation Table: .debug\_abbrev

```

a1: 1
DW_TAG_compile_unit
DW_CHILDREN_yes
DW_AT_name          DW_FORM_string
DW_AT_producer      DW_FORM_string
DW_AT_comp_dir      DW_FORM_string
DW_AT_language_name DW_FORM_data1
DW_AT_low_pc        DW_FORM_addr
DW_AT_high_pc       DW_FORM_data1
DW_AT_stmt_list     DW_FORM_indirect
0
-----
2
DW_TAG_base_type
DW_CHILDREN_no
DW_AT_name          DW_FORM_string
DW_AT_encoding      DW_FORM_data1
DW_AT_byte_size     DW_FORM_data1
0
-----
3
DW_TAG_pointer_type
DW_CHILDREN_no
DW_AT_type          DW_FORM_ref4
0
-----
4
DW_TAG_typedef
DW_CHILDREN_no
DW_AT_name          DW_FORM_string
DW_AT_type          DW_FORM_ref_addr
0
-----
0

```

Figure D.1: Compilation units and abbreviations table

1 **D.1.2 DWARF Stack Operation Examples**

2 *The stack operations defined in Section 2.5.1.3 on page 29. are fairly conventional, but*  
 3 *the following examples illustrate their behavior graphically.*

Before		Operation	After	
0	17	DW_OP_dup	0	17
1	29		1	17
2	1000		2	29
			3	1000
0	17	DW_OP_drop	0	29
1	29		1	1000
2	1000			
0	17	DW_OP_pick, 2	0	1000
1	29		1	17
2	1000		2	29
			3	1000
0	17	DW_OP_over	0	29
1	29		1	17
2	1000		2	29
			3	1000
0	17	DW_OP_swap	0	29
1	29		1	17
2	1000		2	1000
0	17	DW_OP_rot	0	29
1	29		1	1000
2	1000		2	17

### D.1.3 DWARF Location Description Examples

Following are examples of DWARF operations used to form location descriptions:

`DW_OP_reg3`

The value is in register 3.

`DW_OP_regx` (54)

The value is in register 54.

`DW_OP_addr` (0x80d0045c)

The value of a static variable is at machine address 0x80d0045c.

`DW_OP_breg11` (44)

Add 44 to the value in register 11 to get the address of an automatic variable instance.

`DW_OP_fbreg` (-50)

Given a `DW_AT_frame_base` value of “`DW_OP_breg31` 64,” this example computes the address of a local variable that is -50 bytes from a logical frame pointer that is computed by adding 64 to the current stack pointer (register 31).

`DW_OP_bregx` (54, 32)

`DW_OP_deref`

A call-by-reference parameter whose address is in the location beginning 32 bytes from where register 54 points.

`DW_OP_plus_uconst` (4)

A structure member is four bytes from the start of the structure instance. The base address is assumed to be already on the stack.

## Appendix D. Examples (Informative)

1 DW\_OP\_reg3  
2 DW\_OP\_piece (4)  
3 DW\_OP\_reg10  
4 DW\_OP\_piece (2)

5 A variable whose first four bytes reside in register 3 and whose next two  
6 bytes reside in register 10.

7 DW\_OP\_reg0  
8 DW\_OP\_piece (4)  
9 DW\_OP\_piece (4)  
10 DW\_OP\_fpreg (-12)  
11 DW\_OP\_piece (4)

12 A twelve byte value whose first four bytes reside in register zero, whose  
13 middle four bytes are unavailable (perhaps due to optimization), and  
14 whose last four bytes are in memory, 12 bytes before the frame base.

15 DW\_OP\_breg1 (0)  
16 DW\_OP\_breg2 (0)  
17 DW\_OP\_plus  
18 DW\_OP\_stack\_value

19 Add the contents of r1 and r2 to compute a value. This value is the  
20 “contents” of an otherwise anonymous location.

21 DW\_OP\_lit1  
22 DW\_OP\_stack\_value  
23 DW\_OP\_piece (4)  
24 DW\_OP\_breg3 (0)  
25 DW\_OP\_breg4 (0)  
26 DW\_OP\_plus  
27 DW\_OP\_stack\_value  
28 DW\_OP\_piece (4)

29 The object value is found in an anonymous (virtual) location whose value  
30 consists of two parts, given in memory address order: the 4 byte value 1  
31 followed by the four byte value computed from the sum of the contents of  
32 r3 and r4.

## Appendix D. Examples (Informative)

1 `DW_OP_entry_value` (2, `DW_OP_breg1` 0)  
2                   ! *The first operand gives the number of bytes in the*  
3                   ! *second operand (see Section 2.5.1.7 on page 37).*

4           The variable's address is the value that register 1 contained upon entering  
5           the current subprogram.

6 `DW_OP_entry_value` (1, `DW_OP_reg1`)

7           Same as the previous example but uses the more compact register location  
8           description as an operand.

9 `DW_OP_entry_value` (2, `DW_OP_breg1` 0)  
10 `DW_OP_stack_value`

11           The variables's value is the value that register 1 contained upon entering the  
12           current subprogram. This value is the "contents" of an otherwise  
13           anonymous location.

14 `DW_OP_entry_value` (1, `DW_OP_reg1`)  
15 `DW_OP_stack_value`

16           Same as the previous example, but uses the more compact register location  
17           description.

18 `DW_OP_entry_value` (3, `DW_OP_breg4` 16 `DW_OP_deref`)  
19 `DW_OP_stack_value`

20           Add 16 to the value register 4 had upon entering the current subprogram to  
21           form an address and then push the value of the memory location at that  
22           address. This value is the "contents" of an otherwise anonymous location.

23 `DW_OP_entry_value` (1, `DW_OP_reg5`)  
24 `DW_OP_plus_uconst` (16)

25           The address of the memory location is calculated by adding 16 to the value  
26           contained in register 5 upon entering the current subprogram.

## Appendix D. Examples (Informative)

```
1 DW_OP_reg0
2 DW_OP_bit_piece (1, 31)
3 DW_OP_bit_piece (7, 0)
4 DW_OP_reg1
5 DW_OP_piece (1)
```

6 A variable whose first bit resides in the 31st bit of register 0, whose next  
7 seven bits are undefined and whose second byte resides in register 1.

## 8 D.2 Aggregate Examples

9 The following examples illustrate how to represent some of the more  
10 complicated forms of array and record aggregates using DWARF.

### 11 D.2.1 Fortran Simple Array Example

12 Consider the Fortran array source fragment in Figure D.2 following.

```
TYPE array_ptr
REAL :: myvar
REAL, DIMENSION (:), POINTER :: ap
END TYPE array_ptr
TYPE(array_ptr), ALLOCATABLE, DIMENSION(:) :: arrayvar
ALLOCATE(arrayvar(20))
DO I = 1, 20
    ALLOCATE(arrayvar(i)%ap(i+10))
END DO
```

Figure D.2: Fortran array example: source fragment

13 For allocatable and pointer arrays, it is essentially required by the Fortran array  
14 semantics that each array consist of two parts, which we here call 1) the  
15 descriptor and 2) the raw data. (A descriptor has often been called a dope vector  
16 in other contexts, although it is often a structure of some kind rather than a  
17 simple vector.) Because there are two parts, and because the lifetime of the  
18 descriptor is necessarily longer than and includes that of the raw data, there must  
19 be an address somewhere in the descriptor that points to the raw data when, in  
20 fact, there is some (that is, when the “variable” is allocated or associated).

21 For concreteness, suppose that a descriptor looks something like the C structure  
22 in Figure D.3 on the following page. Note, however, that it is a property of the  
23 design that 1) a debugger needs no builtin knowledge of this structure and 2)

## Appendix D. Examples (Informative)

1 there does not need to be an explicit representation of this structure in the  
2 DWARF input to the debugger.

```
struct desc {
    long el_len;           // Element length
    void * base;          // Address of raw data
    int ptr_assoc : 1;    // Pointer is associated flag
    int ptr_alloc : 1;    // Pointer is allocated flag
    int num_dims : 6;     // Number of dimensions
    struct dims_str {    // For each dimension...
        long low_bound;
        long upper_bound;
        long stride;
    } dims[63];
};
```

Figure D.3: Fortran array example: descriptor representation

3 In practice, of course, a “real” descriptor will have dimension substructures only  
4 for as many dimensions as are specified in the num\_dims component. Let us use  
5 the notation desc<n> to indicate a specialization of the desc struct in which n is  
6 the bound for the dims component as well as the contents of the num\_dims  
7 component.

8 Because the arrays considered here come in two parts, it is necessary to  
9 distinguish the parts carefully. In particular, the “address of the variable” or  
10 equivalently, the “base address of the object” *always* refers to the descriptor. For  
11 arrays that do not come in two parts, an implementation can provide a descriptor  
12 anyway, thereby giving it two parts. (This may be convenient for general runtime  
13 support unrelated to debugging.) In this case the above vocabulary applies as  
14 stated. Alternatively, an implementation can do without a descriptor, in which  
15 case the “address of the variable,” or equivalently the “base address of the  
16 object”, refers to the “raw data” (the real data, the only thing around that can be  
17 the object).

18 If an object has a descriptor, then the DWARF type for that object will have a  
19 [DW\\_AT\\_data\\_location](#) attribute. If an object does not have a descriptor, then  
20 usually the DWARF type for the object will not have a [DW\\_AT\\_data\\_location](#)  
21 attribute. (See the following Ada example for a case where the type for an object  
22 without a descriptor does have a [DW\\_AT\\_data\\_location](#) attribute. In that case  
23 the object doubles as its own descriptor.)

## Appendix D. Examples (Informative)

1 The Fortran derived type `array_ptr` can now be re-described in C-like terms that  
2 expose some of the representation as in

```
struct array_ptr {  
    float myvar;  
    desc<1> ap;  
};
```

3 Similarly for variable `arrayvar`:

```
desc<1> arrayvar;
```

4 *Recall that `desc<1>` indicates the 1-dimensional version of `desc`.*

5 Finally, the following notation is useful:

- 6 1. `sizeof(type)`: size in bytes of entities of the given type
- 7 2. `offset(type, comp)`: offset in bytes of the `comp` component within an entity of  
8 the given type

9 The DWARF description is shown in Figure [D.4 on page 308](#).

10 Suppose the program is stopped immediately following completion of the `do`  
11 loop. Suppose further that the user enters the following debug command:

```
debug> print arrayvar(5)%ap(2)
```

12 Interpretation of this expression proceeds as follows:

- 13 1. Lookup name `arrayvar`. We find that it is a variable, whose type is given by  
14 the unnamed type at 6\$. Notice that the type is an array type.
- 15 2. Find the 5<sup>th</sup> element of that array object. To do array indexing requires  
16 several pieces of information:
  - 17 a) the address of the array data
  - 18 b) the lower bounds of the array  
19 [To check that 5 is within bounds would require the upper bound too, but  
20 we will skip that for this example. ]
  - 21 c) the stride

22 For a), check for a [DW\\_AT\\_data\\_location](#) attribute. Since there is one, go  
23 execute the expression, whose result is the address needed. The object  
24 address used in this case is the object we are working on, namely the variable  
25 named `arrayvar`, whose address was found in step 1. (Had there been no



## Appendix D. Examples (Informative)

1 [DW\\_AT\\_data\\_location](#) attribute, the desired address would be the same as  
2 the address from step 1.)

3 For b), for each dimension of the array (only one in this case), go interpret the  
4 usual lower bound attribute. Again this is an expression, which again begins  
5 with [DW\\_OP\\_push\\_object\\_address](#). This object is **still** arrayvar, from step 1,  
6 because we have not begun to actually perform any indexing yet.

7 For c), the default stride applies. Since there is no [DW\\_AT\\_byte\\_stride](#)  
8 attribute, use the size of the array element type, which is the size of type  
9 `array_ptr` (at 3\$).

## Appendix D. Examples (Informative)

part 1 of 2

```
! Description for type of 'ap'
!
1$: DW_TAG_array_type
    ! No name, default (Fortran) ordering, default stride
    DW_AT_type(reference to REAL)
    DW_AT_associated(expression=    ! Test 'ptr_assoc' flag
        DW_OP_push_object_address
        DW_OP_lit<n>                ! where n == offset(ptr_assoc)
        DW_OP_plus
        DW_OP_deref
        DW_OP_lit1                  ! mask for 'ptr_assoc' flag
        DW_OP_and)
    DW_AT_data_location(expression= ! Get raw data address
        DW_OP_push_object_address
        DW_OP_lit<n>                ! where n == offset(base)
        DW_OP_plus
        DW_OP_deref)                ! Type of index of array 'ap'
2$: DW_TAG_subrange_type
    ! No name, default stride
    DW_AT_type(reference to INTEGER)
    DW_AT_lower_bound(expression=
        DW_OP_push_object_address
        DW_OP_lit<n>                ! where n ==
            ! offset(desc, dims) +
            ! offset(dims_str, lower_bound)

        DW_OP_plus
        DW_OP_deref)
    DW_AT_upper_bound(expression=
        DW_OP_push_object_address
        DW_OP_lit<n>                ! where n ==
            ! offset(desc, dims) +
            ! offset(dims_str, upper_bound)

        DW_OP_plus
        DW_OP_deref)
    ! Note: for the m'th dimension, the second operator becomes
    ! DW_OP_lit<n> where
    !     n == offset(desc, dims)          +
    !           (m-1)*sizeof(dims_str)    +
    !           offset(dims_str, [lower|upper]_bound)
    ! That is, the expression does not get longer for each successive
    ! dimension (other than to express the larger offsets involved).
```

Figure D.4: Fortran array example: DWARF description

## Appendix D. Examples (Informative)

part 2 of 2

```
3$: DW_TAG_structure_type
    DW_AT_name("array_ptr")
    DW_AT_byte_size(constant sizeof(REAL) + sizeof(desc<1>))
4$: DW_TAG_member
    DW_AT_name("myvar")
    DW_AT_type(reference to REAL)
    DW_AT_data_member_location(constant 0)
5$: DW_TAG_member
    DW_AT_name("ap");
    DW_AT_type(reference to 1$)
    DW_AT_data_member_location(constant sizeof(REAL))
6$: DW_TAG_array_type
    ! No name, default (Fortran) ordering, default stride
    DW_AT_type(reference to 3$)
    DW_AT_allocated(expression=      ! Test 'ptr_alloc' flag
        DW_OP_push_object_address
        DW_OP_lit<n>                  ! where n == offset(ptr_alloc)
        DW_OP_plus
        DW_OP_deref
        DW_OP_lit2                    ! Mask for 'ptr_alloc' flag
        DW_OP_and)
    DW_AT_data_location(expression=  ! Get raw data address
        DW_OP_push_object_address
        DW_OP_lit<n>                  ! where n == offset(base)
        DW_OP_plus
        DW_OP_deref)
7$: DW_TAG_subrange_type
    ! No name, default stride
    DW_AT_type(reference to INTEGER)
    DW_AT_lower_bound(expression=
        DW_OP_push_object_address
        DW_OP_lit<n>                  ! where n == ...
        DW_OP_plus
        DW_OP_deref)
    DW_AT_upper_bound(expression=
        DW_OP_push_object_address
        DW_OP_lit<n>                  ! where n == ...
        DW_OP_plus
        DW_OP_deref)
8$: DW_TAG_variable
    DW_AT_name("arrayvar")
    DW_AT_type(reference to 6$)
    DW_AT_location(expression=
        ...as appropriate...)      ! Assume static allocation
```

Figure D.4: Fortran array example: DWARF description (*concluded*)

## Appendix D. Examples (Informative)

1 Having acquired all the necessary data, perform the indexing operation in the  
2 usual manner—which has nothing to do with any of the attributes involved up  
3 to now. Those just provide the actual values used in the indexing step.

4 The result is an object within the memory that was dynamically allocated for  
5 arrayvar.

- 6 3. Find the ap component of the object just identified, whose type is array\_ptr.

7 This is a conventional record component lookup and interpretation. It  
8 happens that the ap component in this case begins at offset 4 from the  
9 beginning of the containing object. Component ap has the unnamed array  
10 type defined at 1\$ in the symbol table.

- 11 4. Find the second element of the array object found in step 3. To do array  
12 indexing requires several pieces of information:

- 13 a) the address of the array storage  
14 b) the lower bounds of the array  
15 [To check that 2 is within bounds we would require the upper bound too,  
16 but we will skip that for this example ]  
17 c) the stride

18 This is just like step 2), so the details are omitted. Recall that because the DWARF  
19 type 1\$ has a [DW\\_AT\\_data\\_location](#), the address that results from step 4) is that  
20 of a descriptor, and that address is the address pushed by the  
21 [DW\\_OP\\_push\\_object\\_address](#) operations in 1\$ and 2\$.

22 Note: we happen to be accessing a pointer array here instead of an allocatable  
23 array; but because there is a common underlying representation, the mechanics  
24 are the same. There could be completely different descriptor arrangements and  
25 the mechanics would still be the same—only the stack machines would be  
26 different.

### 27 **D.2.2 Fortran Coarray Examples**

#### 28 **D.2.2.1 Fortran Scalar Coarray Example**

29 The Fortran scalar coarray example in Figure [D.5 on the following page](#) can be  
30 described as illustrated in Figure [D.6 on the next page](#).

## Appendix D. Examples (Informative)

```
INTEGER x[*]
```

Figure D.5: Fortran scalar coarray: source fragment

```
10$: DW_TAG_coarray_type
    DW_AT_type(reference to INTEGER)
    DW_TAG_subrange_type           ! Note omitted upper bound
    DW_AT_lower_bound(constant 1)  ! Can be omitted (default is 1)

11$: DW_TAG_variable
    DW_AT_name("x")
    DW_AT_type(reference to coarray type at 10$)
```

Figure D.6: Fortran scalar coarray: DWARF description

### 1 D.2.2.2 Fortran Array Coarray Example

2 The Fortran (simple) array coarray example in Figure D.7 can be described as  
3 illustrated in Figure D.8.

```
INTEGER x(10)[*]
```

Figure D.7: Fortran array coarray: source fragment

```
10$: DW_TAG_array_type
    DW_AT_ordering(DW_ORD_col_major)
    DW_AT_type(reference to INTEGER)

11$: DW_TAG_subrange_type
    ! DW_AT_lower_bound(constant 1) ! Omitted (default is 1)
    DW_AT_upper_bound(constant 10)

12$: DW_TAG_coarray_type
    DW_AT_type(reference to array type at 10$)

13$: DW_TAG_subrange_type           ! Note omitted upper & lower bounds

14$: DW_TAG_variable
    DW_AT_name("x")
    DW_AT_type(reference to coarray type at 12$)
```

Figure D.8: Fortran array coarray: DWARF description

## Appendix D. Examples (Informative)

### 1 D.2.2.3 Fortran Multidimensional Coarray Example

2 The Fortran multidimensional coarray of a multidimensional array example in  
3 Figure D.9 can be described as illustrated in Figure D.10 following.

```
INTEGER x(10,11,12)[2,3,*]
```

Figure D.9: Fortran multidimensional coarray: source fragment

```
10$: DW_TAG_array_type          ! Note omitted lower bounds (default to 1)
    DW_AT_ordering(DW_ORD_col_major)
    DW_AT_type(reference to INTEGER)
11$: DW_TAG_subrange_type
    DW_AT_upper_bound(constant 10)
12$: DW_TAG_subrange_type
    DW_AT_upper_bound(constant 11)
13$: DW_TAG_subrange_type
    DW_AT_upper_bound(constant 12)

14$: DW_TAG_coarray_type        ! Note omitted lower bounds (default to 1)
    DW_AT_type(reference to array_type at 10$)
15$: DW_TAG_subrange_type
    DW_AT_upper_bound(constant 2)
16$: DW_TAG_subrange_type
    DW_AT_upper_bound(constant 3)
17$: DW_TAG_subrange_type      ! Note omitted upper (& lower) bound

18$: DW_TAG_variable
    DW_AT_name("x")
    DW_AT_type(reference to coarray type at 14$)
```

Figure D.10: Fortran multidimensional coarray: DWARF description

### 1 **D.2.3 Fortran 2008 Assumed-rank Array Example**

2 Consider the example in Figure [D.11](#), which shows an assumed-rank array in  
 3 Fortran 2008 with supplement 29113:<sup>1</sup>

```

SUBROUTINE Foo(x)
  REAL :: x(..)

  ! x has n dimensions

END SUBROUTINE
    
```

Figure D.11: Declaration of a Fortran 2008 assumed-rank array

4 Let's assume the Fortran compiler used an array descriptor that (in C) looks like  
 5 the one shown in Figure [D.12](#).

```

struct array_descriptor {
  void *base_addr;
  int rank;
  struct dim dims[];
}

struct dim {
  int lower_bound;
  int upper_bound;
  int stride;
  int flags;
}
    
```

Figure D.12: One of many possible layouts for an array descriptor

6 The DWARF type for the array *x* can be described as shown in Figure [D.13 on the](#)  
 7 [next page](#).

8 The layout of the array descriptor is not specified by the Fortran standard unless  
 9 the array is explicitly marked as C-interoperable. To get the bounds of an  
 10 assumed-rank array, the expressions in the [DW\\_TAG\\_generic\\_subrange](#) entry  
 11 need to be evaluated for each of the [DW\\_AT\\_rank](#) dimensions as shown by the  
 12 pseudocode in Figure [D.14 on page 315](#).

---

<sup>1</sup>Technical Specification ISO/IEC TS 29113:2012 *Further Interoperability of Fortran with C*

## Appendix D. Examples (Informative)

```

10$: DW_TAG_array_type
    DW_AT_type(reference to real)
    DW_AT_rank(expression=
        DW_OP_push_object_address
        DW_OP_lit<n>                ! offset of rank in descriptor
        DW_OP_plus
        DW_OP_deref)
    DW_AT_data_location(expression=
        DW_OP_push_object_address
        DW_OP_lit<n>                ! offset of data in descriptor
        DW_OP_plus
        DW_OP_deref)
11$: DW_TAG_generic_subrange
    DW_AT_type(reference to integer)
    DW_AT_lower_bound(expression=
        ! Looks up the lower bound of dimension i.
        ! Operation                ! Stack effect
        ! (implicit)                ! i
        DW_OP_lit<n>                ! i sizeof(dim)
        DW_OP_mul                    ! dim[i]
        DW_OP_lit<n>                ! dim[i] offsetof(dim)
        DW_OP_plus                    ! dim[i]+offset
        DW_OP_push_object_address    ! dim[i]+offsetof(dim) objptr
        DW_OP_plus                    ! objptr.dim[i]
        DW_OP_lit<n>                ! objptr.dim[i] offsetof(lb)
        DW_OP_plus                    ! objptr.dim[i].lowerbound
        DW_OP_deref)                ! *objptr.dim[i].lowerbound
    DW_AT_upper_bound(expression=
        ! Looks up the upper bound of dimension i.
        DW_OP_lit<n>                ! sizeof(dim)
        DW_OP_mul                    !
        DW_OP_lit<n>                ! offsetof(dim)
        DW_OP_plus                    !
        DW_OP_push_object_address    !
        DW_OP_plus                    !
        DW_OP_lit<n>                ! offset of upperbound in dim
        DW_OP_plus                    !
        DW_OP_deref)
    DW_AT_byte_stride(expression=
        ! Looks up the byte stride of dimension i.
        ...
        ! (analogous to DW_AT_upper_bound)
    )

```

Figure D.13: Sample DWARF for the array descriptor in Figure [D.12](#)



## Appendix D. Examples (Informative)

```
typedef struct {
    int lower, upper, stride;
} dims_t;

typedef struct {
    int rank;
    struct dims_t *dims;
} array_t;

array_t get_dynamic_array_dims(DW_TAG_array a) {
    array_t result;

    // Evaluate the DW_AT_rank expression to get the
    // number of dimensions.
    dwarf_stack_t stack;
    dwarf_eval(stack, a.rank_expr);
    result.rank = dwarf_pop(stack);
    result.dims = new dims_t[result.rank];

    // Iterate over all dimensions and find their bounds.
    for (int i = 0; i < result.rank; i++) {
        // Evaluate the generic subrange's DW_AT_lower
        // expression for dimension i.
        dwarf_push(stack, i);
        assert( stack.size == 1 );
        dwarf_eval(stack, a.generic_subrange.lower_expr);
        result.dims[i].lower = dwarf_pop(stack);
        assert( stack.size == 0 );

        dwarf_push(stack, i);
        dwarf_eval(stack, a.generic_subrange.upper_expr);
        result.dims[i].upper = dwarf_pop(stack);

        dwarf_push(stack, i);
        dwarf_eval(stack, a.generic_subrange.byte_stride_expr);
        result.dims[i].stride = dwarf_pop(stack);
    }
    return result;
}
```

Figure D.14: How to interpret the DWARF from Figure D.13

## D.2.4 Fortran Dynamic Type Example

Consider the Fortran 90 example of dynamic properties in Figure D.15. This can be represented in DWARF as illustrated in Figure D.16 on the following page.

Note that unnamed dynamic types are used to avoid replicating the full description of the underlying type `dt` that is shared by several variables.

```

PROGRAM Sample

  TYPE :: dt (1)
    INTEGER, LEN :: 1
    INTEGER :: arr(1)
  END TYPE

  INTEGER :: n = 4
  CONTAINS

  SUBROUTINE S()
    TYPE (dt(n))           :: t1
    TYPE (dt(n)), pointer  :: t2
    TYPE (dt(n)), allocatable :: t3, t4
  END SUBROUTINE

END Sample

```

Figure D.15: Fortran dynamic type example: source

## Appendix D. Examples (Informative)

```
11$: DW_TAG_structure_type
      DW_AT_name("dt")
      DW_TAG_member
          ...
...
13$: DW_TAG_dynamic_type          ! plain version
      DW_AT_data_location (dwarf expression to locate raw data)
      DW_AT_type (11$)
14$: DW_TAG_dynamic_type          ! 'pointer' version
      DW_AT_data_location (dwarf expression to locate raw data)
      DW_AT_associated (dwarf expression to test if associated)
      DW_AT_type (11$)
15$: DW_TAG_dynamic_type          ! 'allocatable' version
      DW_AT_data_location (dwarf expression to locate raw data)
      DW_AT_allocated (dwarf expression to test is allocated)
      DW_AT_type (11$)
16$: DW_TAG_variable
      DW_AT_name ("t1")
      DW_AT_type (13$)
      DW_AT_location (dwarf expression to locate descriptor)
17$: DW_TAG_variable
      DW_AT_name ("t2")
      DW_AT_type (14$)
      DW_AT_location (dwarf expression to locate descriptor)
18$: DW_TAG_variable
      DW_AT_name ("t3")
      DW_AT_type (15$)
      DW_AT_location (dwarf expression to locate descriptor)
19$: DW_TAG_variable
      DW_AT_name ("t4")
      DW_AT_type (15$)
      DW_AT_location (dwarf expression to locate descriptor)
```

Figure D.16: Fortran dynamic type example: DWARF description

## 1 D.2.5 C/C++ Anonymous Structure Example

2 An example of a C/C++ structure is shown in Figure D.17. For this source, the  
 3 DWARF description in Figure D.18 is appropriate. In this example, b is  
 4 referenced as if it were defined in the enclosing structure foo.

```

struct foo {
    int a;
    struct {
        int b;
    };
} x;

void bar(void)
{
    struct foo t;
    t.a = 1;
    t.b = 2;
}

```

Figure D.17: Anonymous structure example: source fragment

```

1$: DW_TAG_structure_type
    DW_AT_name("foo")
2$: DW_TAG_member
    DW_AT_name("a")
3$: DW_TAG_structure_type
    DW_AT_export_symbols
4$: DW_TAG_member
    DW_AT_name("b")

```

Figure D.18: Anonymous structure example: DWARF description

## 5 D.2.6 Ada Example

6 Figure D.19 on the next page illustrates two kinds of Ada parameterized array,  
 7 one embedded in a record.

8 VEC1 illustrates an (unnamed) array type where the upper bound of the first and  
 9 only dimension is determined at runtime. Ada semantics require that the value  
 10 of an array bound is fixed at the time the array type is elaborated (where  
 11 *elaboration* refers to the runtime executable aspects of type processing). For the  
 12 purposes of this example, we assume that there are no other assignments to M so  
 13 that it safe for the REC1 type description to refer directly to that variable (rather  
 14 than a compiler-generated copy).

## Appendix D. Examples (Informative)

```
M : INTEGER := <exp>;
VEC1 : array (1..M) of INTEGER;
subtype TEENY is INTEGER range 1..100;
type ARR is array (INTEGER range <>) of INTEGER;
type REC2(N : TEENY := 100) is record
    VEC2 : ARR(1..N);
end record;

OBJ2B : REC2;
```

Figure D.19: Ada example: source fragment

1 REC2 illustrates another array type (the unnamed type of component VEC2) where  
2 the upper bound of the first and only bound is also determined at runtime. In  
3 this case, the upper bound is contained in a discriminant of the containing record  
4 type. (A *discriminant* is a component of a record whose value cannot be changed  
5 independently of the rest of the record because that value is potentially used in  
6 the specification of other components of the record.)

7 The DWARF description is shown in Figure [D.20 on the following page](#).

8 Interesting aspects about this example are:

- 9 1. The array VEC2 is “immediately” contained within structure REC2 (there is no  
10 intermediate descriptor or indirection), which is reflected in the absence of a  
11 [DW\\_AT\\_data\\_location](#) attribute on the array type at 28\$.
- 12 2. One of the bounds of VEC2 is nonetheless dynamic and part of the same  
13 containing record. It is described as a reference to a member, and the location  
14 of the upper bound is determined as for any member. That is, the location is  
15 determined using an address calculation relative to the base of the containing  
16 object.  
17 A consumer must notice that the referenced bound is a member of the same  
18 containing object and implicitly push the base address of the containing  
19 object just as for accessing a data member generally.
- 20 3. The lack of a subtype concept in DWARF means that DWARF types serve the  
21 role of subtypes and must replicate information from the parent type. For this  
22 reason, DWARF for the unconstrained array type ARR is not needed for the  
23 purposes of this example and therefore is not shown.

## Appendix D. Examples (Informative)

```

11$: DW_TAG_variable
    DW_AT_name("M")
    DW_AT_type(reference to INTEGER)
12$: DW_TAG_array_type
    ! No name, default (Ada) order, default stride
    DW_AT_type(reference to INTEGER)
13$: DW_TAG_subrange_type
    DW_AT_type(reference to INTEGER)
    DW_AT_lower_bound(constant 1)
    DW_AT_upper_bound(reference to variable M at 11$)
14$: DW_TAG_variable
    DW_AT_name("VEC1")
    DW_AT_type(reference to array type at 12$)
    . . .
21$: DW_TAG_subrange_type
    DW_AT_name("TEENY")
    DW_AT_type(reference to INTEGER)
    DW_AT_lower_bound(constant 1)
    DW_AT_upper_bound(constant 100)
    . . .
26$: DW_TAG_structure_type
    DW_AT_name("REC2")
27$: DW_TAG_member
    DW_AT_name("N")
    DW_AT_type(reference to subtype TEENY at 21$)
    DW_AT_data_member_location(constant 0)
28$: DW_TAG_array_type
    ! No name, default (Ada) order, default stride
    ! Default data location
    DW_AT_type(reference to INTEGER)
29$: DW_TAG_subrange_type
    DW_AT_type(reference to subrange TEENY at 21$)
    DW_AT_lower_bound(constant 1)
    DW_AT_upper_bound(reference to member N at 27$)
30$: DW_TAG_member
    DW_AT_name("VEC2")
    DW_AT_type(reference to array "subtype" at 28$)
    DW_AT_data_member_location(machine=
        DW_OP_lit<n>           ! where n == offset(REC2, VEC2)
        DW_OP_plus)
    . . .
41$: DW_TAG_variable
    DW_AT_name("OBJ2B")
    DW_AT_type(reference to REC2 at 26$)
    DW_AT_location(...as appropriate...)

```

Figure D.20: Ada example: DWARF description

## Appendix D. Examples (Informative)

### D.2.7 Pascal Example

The Pascal source in Figure D.21 following is used to illustrate the representation of packed unaligned bit fields.

```
TYPE T : PACKED RECORD                                { bit size is 2  }
      F5 : BOOLEAN;                                  { bit offset is 0 }
      F6 : BOOLEAN;                                  { bit offset is 1 }
      END;
VAR V : PACKED RECORD
      F1 : BOOLEAN;                                  { bit offset is 0 }
      F2 : PACKED RECORD                             { bit offset is 1 }
          F3 : INTEGER;                              { bit offset is 0 in F2,
                                                    1 in V }
      END;
      F4 : PACKED ARRAY [0..1] OF T; { bit offset is 33 }
      F7 : T;                                         { bit offset is 37 }
      END;
```

Figure D.21: Packed record example: source fragment

The DWARF representation in Figure D.22 is appropriate. `DW_TAG_packed_type` entries could be added to better represent the source, but these do not otherwise affect the example and are omitted for clarity. Note that this same representation applies to both typical big- and little-endian architectures using the conventions described in Section 5.7.6 on page 122.

*part 1 of 2*

```
10$: DW_TAG_base_type
      DW_AT_name("BOOLEAN")
      ...
11$: DW_TAG_base_type
      DW_AT_name("INTEGER")
      ...
20$: DW_TAG_structure_type
      DW_AT_name("T")
      DW_AT_bit_size(2)
      DW_TAG_member
          DW_AT_name("F5")
          DW_AT_type(reference to 10$)
          DW_AT_data_bit_offset(0)      ! may be omitted
          DW_AT_bit_size(1)
```

Figure D.22: Packed record example: DWARF description

## Appendix D. Examples (Informative)

part 2 of 2

```

    DW_TAG_member
      DW_AT_name("F6")
      DW_AT_type(reference to 10$)
      DW_AT_data_bit_offset(1)
      DW_AT_bit_size(1)
21$: DW_TAG_structure_type           ! anonymous type for F2
      DW_TAG_member
        DW_AT_name("F3")
        DW_AT_type(reference to 11$)
22$: DW_TAG_array_type             ! anonymous type for F4
      DW_AT_type(reference to 20$)
      DW_TAG_subrange_type
        DW_AT_type(reference to 11$)
        DW_AT_lower_bound(0)
        DW_AT_upper_bound(1)
      DW_AT_bit_stride(2)
      DW_AT_bit_size(4)
23$: DW_TAG_structure_type           ! anonymous type for V
      DW_AT_bit_size(39)
      DW_TAG_member
        DW_AT_name("F1")
        DW_AT_type(reference to 10$)
        DW_AT_data_bit_offset(0)      ! may be omitted
        DW_AT_bit_size(1) ! may be omitted
      DW_TAG_member
        DW_AT_name("F2")
        DW_AT_type(reference to 21$)
        DW_AT_data_bit_offset(1)
        DW_AT_bit_size(32) ! may be omitted
      DW_TAG_member
        DW_AT_name("F4")
        DW_AT_type(reference to 22$)
        DW_AT_data_bit_offset(33)
        DW_AT_bit_size(4) ! may be omitted
      DW_TAG_member
        DW_AT_name("F7")
        DW_AT_type(reference to 20$)   ! type T
        DW_AT_data_bit_offset(37)
        DW_AT_bit_size(2)             ! may be omitted
      DW_TAG_variable
        DW_AT_name("V")
        DW_AT_type(reference to 23$)
        DW_AT_location(...)
    ...

```

Figure D.22: Packed record example: DWARF description (*concluded*)



## 1 D.2.8 C/C++ Bit-Field Examples

2 *Bit fields in C and C++ typically require the use of the [DW\\_AT\\_data\\_bit\\_offset](#) and*  
 3 *[DW\\_AT\\_bit\\_size](#) attributes.*

4 *This Standard uses the following bit numbering and direction conventions in examples.*  
 5 *These conventions are for illustrative purposes and other conventions may apply on*  
 6 *particular architectures.*

- 7 • *For big-endian architectures, bit offsets are counted from high-order to low-order*  
 8 *bits within a byte (or larger storage unit); in this case, the bit offset identifies the*  
 9 *high-order bit of the object.*
- 10 • *For little-endian architectures, bit offsets are counted from low-order to high-order*  
 11 *bits within a byte (or larger storage unit); in this case, the bit offset identifies the*  
 12 *low-order bit of the object.*

13 *In either case, the bit so identified is defined as the beginning of the object.*

14 This section illustrates one possible representation of the following C structure  
 15 definition in both big- and little-endian byte orders:

```

struct S {
    int j:5;
    int k:6;
    int m:5;
    int n:8;
};
```

16 Figures [D.23](#) and [D.24](#) on the next page show the structure layout and data bit  
 17 offsets for example big- and little-endian architectures, respectively. Both  
 18 diagrams show a structure that begins at address A and whose size is four bytes.  
 19 Also, high order bits are to the left and low order bits are to the right.

20 Note that data member bit offsets in this example are the same for both big- and  
 21 little-endian architectures even though the fields are allocated in different  
 22 directions (high-order to low-order versus low-order to high-order); the bit  
 23 naming conventions for memory and/or registers of the target architecture may  
 24 or may not make this seem natural.

## Appendix D. Examples (Informative)

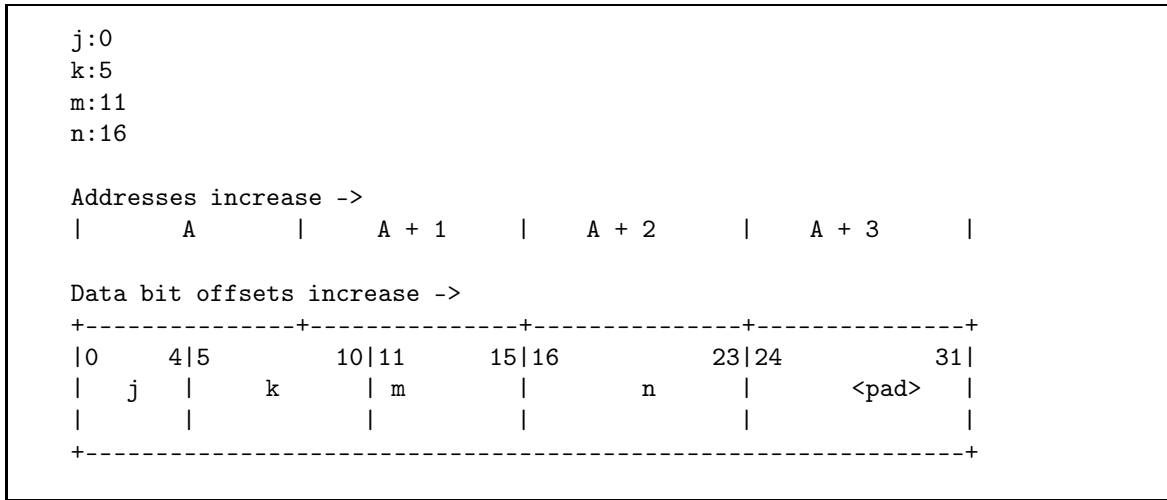


Figure D.23: Big-endian data bit offsets




Figure D.24: Little-endian data bit offsets

### D.2.9 Ada Biased Bit-Field Example

The Ada source in Figure D.25 on the following page demonstrates how a member of a record, which normally occupies six bits, can be biased to fit into three bits when the range is known. The encoded values [ 0..7 ] correspond to the source values [ 50..57 ] used by the application.

The DWARF description is shown in Figure D.26 on the next page. The bias chosen, which in this case corresponds to the lower bound, is specified in the base type at 1\$.

## Appendix D. Examples (Informative)

```
type SmallRangeType is range 50 .. 57;
type RecordType is record
  A : SmallRangeType;
end record;
for RecordType use record
  A at 0 range 0 .. 2;
end record;
LocalRecord : RecordType;
```

Figure D.25: Ada biased bit-field example: Ada source

```
1$: DW_TAG_base_type
    DW_AT_byte_size(1)
    DW_AT_encoding(DW_ATE_unsigned)
    DW_AT_bias(50)
    DW_AT_artificial(1)
2$: DW_TAG_subrange_type
    DW_AT_name("SmallRangeType")
    DW_AT_lower_bound(50)
    DW_AT_upper_bound(57)
    DW_AT_type(reference to 1$)
3$: DW_TAG_structure_type
    DW_AT_name("RecordType")
    DW_AT_byte_size(1)
4$: DW_TAG_member
    DW_AT_name("A")
    DW_AT_type(reference to 2$)
    DW_AT_bit_size(3)
    DW_AT_data_bit_offset(0)
5$: DW_TAG_variable
    DW_AT_name("LocalRecord")
    DW_AT_type(reference to 3$)
    DW_AT_location ...
```

Figure D.26: Ada biased bit-field example: DWARF description

- 1 Note that other choices of encoding and bias lead to the same result. For example,
- 2 the `DW_ATE_signed` encoding can be used in combination with a bias of 54.
- 3 If the valid range of values is completely negative (for example, -57..-50) then
- 4 only signed encoding is valid, and the bias will also need to be negative (-53).

## 1 **D.2.10 Variant Entry Examples**

2 The following examples illustrate some of the diverse ways that the DWARF  
3 variant entry constructs are used in various programming languages.

### 4 **D.2.10.1 Pascal Variant Entry Example**

5 A Pascal record example without a variant part is shown in [D.2.7 on page 321](#).  
6 Here a Pascal record with a variant part is shown in [Figure D.27](#) following. The  
7 corresponding DWARF representation follows in [Figure D.28 on the next page](#).

```
RPoint = Record
  Case UsePolar : Boolean of
    False : (X, Y : Real);
    True  : (Radius, Theta : Real);
  end;
end;
```

Figure D.27: Pascal variant record example: source

## Appendix D. Examples (Informative)

```
! Description for type RPoint
!
1$: DW_TAG_structure_type
    DW_AT_name("RPoint")
    DW_TAG_variant_part
        DW_AT_discr (reference to 2$)
2$:    DW_TAG_member
        DW_AT_name("UsePolar")
        DW_AT_type(reference to Boolean)
    DW_TAG_variant
        DW_AT_discr_value(constant 0)
        DW_TAG_member
            DW_AT_name("X")
            DW_AT_type(reference to Real)
            DW_AT_data_member_location(1)
        DW_TAG_member
            DW_AT_name("Y")
            DW_AT_type(reference to Real)
            DW_AT_data_member_location(5)
    DW_TAG_variant
        DW_AT_discr_value(constant 1)
        DW_TAG_member
            DW_AT_name("Radius")
            DW_AT_type(reference to Real)
            DW_AT_data_member_location(1)
        DW_TAG_member
            DW_AT_name("Theta")
            DW_AT_type(reference to Real)
            DW_AT_data_member_location(5)
```

Figure D.28: Pascal variant record example: DWARF description

1 Notice that the "tag" (member UsePolar in this case) is the first child of the  
2 variant part. A "tagless" version of this example would simply delete "UsePolar :"  
3 from the second line of the source (so that the tag has no name, hence is not  
4 visible). In the DWARF description, the member entry and name for UsePolar  
5 are then deleted, as is the `DW_AT_discr` attribute, and the remaining type  
6 attribute is made an attribute of the containing variant part entry.

1 **D.2.10.2 Ada Variant Entry Example**

2 An Ada example variant part is illustrated in Figure D.29 following. The  
3 corresponding DWARF is shown in Figure D.30 on the following page.

```
type R (D : integer) is
  record
    A : integer;
    case D is
      when 0 =>
        F : float;
      when 1 =>
        N : integer;
      when others =>
        null;
    end case;
  end record;
```

Figure D.29: Ada variant record example: source

4 For Ada, note that the tag is not "declared" as part of the variant part construct.  
5 Rather the variant part refers to a discriminant of the containing type which  
6 necessarily occurs as an initial member in the sequence of record components.  
7 This reference is implemented as a [DW\\_AT\\_discr](#) attribute of the  
8 [DW\\_TAG\\_variant\\_part](#) entry.

## Appendix D. Examples (Informative)

```
DW_TAG_structure_type
  DW_AT_name("r")
1$: DW_TAG_member          ! Discriminant
    DW_AT_type(reference to integer)
    DW_AT_data_member_location(DW_OP_plus_uconst 0)
    DW_AT_name("d")
  DW_TAG_member
    DW_AT_type(reference to integer)
    DW_AT_data_member_location(DW_OP_plus_uconst 4)
    DW_AT_name("a")
  DW_TAG_variant_part
    DW_AT_discr(reference to 1$)
    DW_TAG_variant
      DW_AT_discr_value(0)
      DW_TAG_member
        DW_AT_type(reference to float)
        DW_AT_data_member_location(DW_OP_plus_uconst 8)
        DW_AT_name("f")
    DW_TAG_variant
      DW_AT_discr_value(1)
      DW_TAG_member
        DW_AT_type(reference to integer)
        DW_AT_data_member_location(DW_OP_plus_uconst 8)
        DW_AT_name("n")
  DW_TAG_variant
    ! No members described for the "others" variant
```

Figure D.30: Ada variant record example: DWARF description

### 1 D.2.10.3 Rust Enum Example

2 While Rust does not have a variant record concept similar to that in Pascal or  
3 Ada, it does use a similar mechanism in the implementation of enums. To  
4 illustrate, consider the enumeration in Figure D.31 following. This can be  
5 described in DWARF as shown in Figure D.32 on the next page.

```
enum Message {
  F(f64),
  U(u32),
  N(i32)
}
```

Figure D.31: Rust enum example: source

## Appendix D. Examples (Informative)

```
DW_TAG_structure_type
  DW_AT_name("Message")
  DW_TAG_variant_part
    DW_AT_discr(reference to $1)
$1:  DW_TAG_member          ! Artificial discriminant
      DW_AT_type(reference to u32)
      DW_AT_data_member_location(0)
      DW_AT_artificial(1)
    DW_TAG_variant
      DW_AT_discr_value(0)
      DW_TAG_member
        DW_AT_type(reference to f32)
        DW_AT_name("F")
        DW_AT_data_member_location(4)
    DW_TAG_variant
      DW_AT_discr_value(1)
      DW_TAG_member
        DW_AT_type(reference to u32)
        DW_AT_name("U")
        DW_AT_data_member_location(4)
    DW_TAG_variant
      DW_AT_discr_value(2)
      DW_TAG_member
        DW_AT_type(reference to i32)
        DW_AT_name("N")
        DW_AT_data_member_location(4)
```

Figure D.32: Rust enum example: DWARF description



## 1 D.3 Namespace Examples

2 The C++ example in Figure D.33 is used to illustrate the representation of  
 3 namespaces. The DWARF representation in Figure D.34 on the next page is  
 4 appropriate.

```

namespace {
    int i;
}
namespace A {
    namespace B {
        int j;
        int myfunc (int a);
        float myfunc (float f) { return f - 2.0; }
        int myfunc2(int a) { return a + 2; }
    }
}
namespace Y {
    using A::B::j;           // (1) using declaration
    int foo;
}
using A::B::j;             // (2) using declaration
namespace Foo = A::B;     // (3) namespace alias
using Foo::myfunc;        // (4) using declaration
using namespace Foo;      // (5) using directive
namespace A {
    namespace B {
        using namespace Y; // (6) using directive
        int k;
    }
}
int Foo::myfunc(int a)
{
    i = 3;
    j = 4;
    return myfunc2(3) + j + i + a + 2;
}

```

Figure D.33: Namespace example #1: source fragment

## Appendix D. Examples (Informative)

part 1 of 2

```
1$: DW_TAG_base_type
    DW_AT_name("int")
    ...
2$: DW_TAG_base_type
    DW_AT_name("float")
    ...
6$: DW_TAG_namespace
    ! no DW_AT_name attribute
    DW_AT_export_symbols          ! Implied by C++, but can be explicit
    DW_TAG_variable
        DW_AT_name("i")
        DW_AT_type(reference to 1$)
        DW_AT_location ...
    ...
10$: DW_TAG_namespace
    DW_AT_name("A")
20$: DW_TAG_namespace
    DW_AT_name("B")
30$: DW_TAG_variable
    DW_AT_name("j")
    DW_AT_type(reference to 1$)
    DW_AT_location ...
    ...
34$: DW_TAG_subprogram
    DW_AT_name("myfunc")
    DW_AT_type(reference to 1$)
    ...
36$: DW_TAG_subprogram
    DW_AT_name("myfunc")
    DW_AT_type(reference to 2$)
    ...
38$: DW_TAG_subprogram
    DW_AT_name("myfunc2")
    DW_AT_low_pc ...
    DW_AT_high_pc ...
    DW_AT_type(reference to 1$)
    ...
```

Figure D.34: Namespace example #1: DWARF description

## Appendix D. Examples (Informative)

part 2 of 2

```
40$: DW_TAG_namespace
    DW_AT_name("Y")
    DW_TAG_imported_declaration      ! (1) using-declaration
        DW_AT_import(reference to 30$)
    DW_TAG_variable
        DW_AT_name("foo")
        DW_AT_type(reference to 1$)
        DW_AT_location ...
    ...
    DW_TAG_imported_declaration      ! (2) using declaration
        DW_AT_import(reference to 30$)
    DW_TAG_imported_declaration      ! (3) namespace alias
        DW_AT_name("Foo")
        DW_AT_import(reference to 20$)
    DW_TAG_imported_declaration      ! (4) using declaration
        DW_AT_import(reference to 34$)      ! - part 1
    DW_TAG_imported_declaration      ! (4) using declaration
        DW_AT_import(reference to 36$)      ! - part 2
    DW_TAG_imported_module            ! (5) using directive
        DW_AT_import(reference to 20$)
    DW_TAG_namespace
        DW_AT_extension(reference to 10$)
        DW_TAG_namespace
            DW_AT_extension(reference to 20$)
            DW_TAG_imported_module      ! (6) using directive
                DW_AT_import(reference to 40$)
            DW_TAG_variable
                DW_AT_name("k")
                DW_AT_type(reference to 1$)
                DW_AT_location ...
        ...
60$: DW_TAG_subprogram
    DW_AT_specification(reference to 34$)
    DW_AT_low_pc ...
    DW_AT_high_pc ...
    ...
```

Figure D.34: Namespace example #1: DWARF description (*concluded*)

## Appendix D. Examples (Informative)

1 As a further namespace example, consider the inlined namespace shown in  
2 Figure D.35. For this source, the DWARF description in Figure D.36 is  
3 appropriate. In this example, `a` may be referenced either as a member of the fully  
4 qualified namespace `A::B`, or as if it were defined in the enclosing namespace, `A`.

```
namespace A {
    inline namespace B { // (1) inline namespace
        int a;
    }
}

void foo (void)
{
    using A::B::a;
    a = 1;
}

void bar (void)
{
    using A::a;
    a = 2;
}
```

Figure D.35: Namespace example #2: source fragment

```
1$: DW_TAG_namespace
    DW_AT_name("A")
2$: DW_TAG_namespace
    DW_AT_name("B")
    DW_AT_export_symbols
3$: DW_TAG_variable
    DW_AT_name("a")
```

Figure D.36: Namespace example #2: DWARF description

## D.4 Member Function Examples

Consider the member function example fragment in Figure D.37. The DWARF representation in Figure D.38 is appropriate.

```
class A
{
    void func1(int x1);
    void func2() const;
    static void func3(int x3);
};
void A::func1(int x) {}
```

Figure D.37: Member function example: source fragment

*part 1 of 2*

```
2$: DW_TAG_base_type
    DW_AT_name("int")
    ...
3$: DW_TAG_class_type
    DW_AT_name("A")
    ...
4$: DW_TAG_pointer_type
    DW_AT_type(reference to 3$)
    ...
5$: DW_TAG_const_type
    DW_AT_type(reference to 3$)
    ...
6$: DW_TAG_pointer_type
    DW_AT_type(reference to 5$)
    ...
7$: DW_TAG_subprogram
    DW_AT_declaration
    DW_AT_name("func1")
    DW_AT_object_pointer(reference to 8$)
        ! References a formal parameter in this
        ! member function
    ...
```

Figure D.38: Member function example: DWARF description

## Appendix D. Examples (Informative)

part 2 of 2

```
8$:      DW_TAG_formal_parameter
        DW_AT_artificial(true)
        DW_AT_name("this")
        DW_AT_type(reference to 4$)
        ! Makes type of 'this' as 'A*' =>
        ! func1 has not been marked const
        ! or volatile
        DW_AT_location ...
        ...
9$:      DW_TAG_formal_parameter
        DW_AT_name(x1)
        DW_AT_type(reference to 2$)
        ...
10$:    DW_TAG_subprogram
        DW_AT_declaration
        DW_AT_name("func2")
        DW_AT_object_pointer(reference to 11$)
        ! References a formal parameter in this
        ! member function
        ...
11$:    DW_TAG_formal_parameter
        DW_AT_artificial(true)
        DW_AT_name("this")
        DW_AT_type(reference to 6$)
        ! Makes type of 'this' as 'A const*' =>
        !   func2 marked as const
        DW_AT_location ...
        ...
12$:    DW_TAG_subprogram
        DW_AT_declaration
        DW_AT_name("func3")
        ...
        ! No object pointer reference formal parameter
        ! implies func3 is static
13$:    DW_TAG_formal_parameter
        DW_AT_name(x3)
        DW_AT_type(reference to 2$)
        ...
```

Figure D.38: Member function example: DWARF description (*concluded*)

## Appendix D. Examples (Informative)

1 As a further example illustrating &- and &&-qualification of member functions,  
2 consider the member function example fragment in Figure D.39. The DWARF  
3 representation in Figure D.40 on the following page is appropriate.

```
class A {
public:
    void f() const &&;
};

void g() {
    A a;
    // The type of pointer is "void (A::*)() const &&".
    auto pointer_to_member_function = &A::f;
}
```

Figure D.39: Reference- and rvalue-reference-qualification example: source fragment

## Appendix D. Examples (Informative)

```
100$: DW_TAG_class_type
      DW_AT_name("A")
      DW_TAG_subprogram
        DW_AT_name("f")
        DW_AT_rvalue_reference(0x01)
        DW_TAG_formal_parameter
          DW_AT_type(ref to 200$)      ! to const A*
          DW_AT_artificial(0x01)

200$: ! const A*
      DW_TAG_pointer_type
        DW_AT_type(ref to 300$)      ! to const A

300$: ! const A
      DW_TAG_const_type
        DW_AT_type(ref to 100$)      ! to class A

400$: ! mfpnr
      DW_TAG_ptr_to_member_type
        DW_AT_type(ref to 500$)      ! to functype
        DW_AT_containing_type(ref to 100$) ! to class A

500$: ! functype
      DW_TAG_subroutine_type
        DW_AT_rvalue_reference(0x01)
        DW_TAG_formal_parameter
          DW_AT_type(ref to 200$)      ! to const A*
          DW_AT_artificial(0x01)

600$: DW_TAG_subprogram
      DW_AT_name("g")
      DW_TAG_variable
        DW_AT_name("a")
        DW_AT_type(ref to 100$)      ! to class A
      DW_TAG_variable
        DW_AT_name("pointer_to_member_function")
        DW_AT_type(ref to 400$)
```

Figure D.40: Reference- and rvalue-reference-qualification example: DWARF description



## D.5 Line Number Examples

### D.5.1 Line Number Header Example

Figure D.41 illustrates a line number header (see Section 6.2.4 on page 160). There are multiple alternative filename formats, which include the source and URL types.

Field Number	Field Name	Value(s)
1	unit_length	<unit length>
2	version	6
3	address_size	4 or 8
4	<i>Reserved</i>	0
5	header_length	<header length>
6	minimum_instruction_length	1
7	maximum_operations_per_instruction	1
8	default_is_stmt	1 (true)
9	line_base	-3
10	line_range	12
11	opcode_base	13
12	standard_opcode_lengths	[0,1,1,1,1,0,0,0,0,0,0,1]
13	directory_format_count	1
14	directory_format_table	[DW_LNCT_path, DW_FORM_string], [0, 0]
15	directories_count	1
16	directories	[0, <directory path string>]
17	file_name_format_count	3
18	file_name_format_table	[DW_LNCT_source, DW_FORM_strp], [0, 0] [DW_LNCT_path, DW_FORM_string], [DW_LNCT_directory_index, DW_FORM_udata], [DW_LNCT_timestamp, DW_FORM_udata], [DW_LNCT_size, DW_FORM_udata], [0, 0], [DW_LNCT_URL, DW_FORM_strp], [0, 0]
19	file_names_count	4
20	file_names	[0, {<source string offset>}], [2, {<URL string offset>}], [1, {<name string 1>, <directory index=0>, <timestamp 1>, <size 1>}], [1, {<name string 2>, <directory index=0>, <timestamp 2>, <size 2>}]

Figure D.41: Example line number program header

## D.5.2 Line Number Special Opcode Example

Given the example header in Figure D.41 on the previous page, we can use a special opcode whenever two successive rows in the matrix have source line numbers differing by any value within the range [-3, 8] and (because of the limited number of opcodes available) when the difference between addresses is within the range [0, 20]. The resulting opcode mapping is shown in Figure D.42.

Note in the bottom row of the figure that not all line advances are available for the maximum operation advance.

Operation Advance	Line Advance											
	-3	-2	-1	0	1	2	3	4	5	6	7	8
0	13	14	15	16	17	18	19	20	21	22	23	24
1	25	26	27	28	29	30	31	32	33	34	35	36
2	37	38	39	40	41	42	43	44	45	46	47	48
3	49	50	51	52	53	54	55	56	57	58	59	60
4	61	62	63	64	65	66	67	68	69	70	71	72
5	73	74	75	76	77	78	79	80	81	82	83	84
6	85	86	87	88	89	90	91	92	93	94	95	96
7	97	98	99	100	101	102	103	104	105	106	107	108
8	109	110	111	112	113	114	115	116	117	118	119	120
9	121	122	123	124	125	126	127	128	129	130	131	132
10	133	134	135	136	137	138	139	140	141	142	143	144
11	145	146	147	148	149	150	151	152	153	154	155	156
12	157	158	159	160	161	162	163	164	165	166	167	168
13	169	170	171	172	173	174	175	176	177	178	179	180
14	181	182	183	184	185	186	187	188	189	190	191	192
15	193	194	195	196	197	198	199	200	201	202	203	204
16	205	206	207	208	209	210	211	212	213	214	215	216
17	217	218	219	220	221	222	223	224	225	226	227	228
18	229	230	231	232	233	234	235	236	237	238	239	240
19	241	242	243	244	245	246	247	248	249	250	251	252
20	253	254	255									

Figure D.42: Example line number special opcode mapping

There is no requirement that the expression  $255 - \text{line\_base} + 1$  be an integral multiple of  $\text{line\_range}$ .

### D.5.3 Line Number Program Example

Consider the simple source file and the resulting machine code for the Intel 8086 processor in Figure D.43.

```

1: int
2: main()
   0x239: push pb
   0x23a: mov bp,sp
3: {
4: printf("Omit needless words\n");
   0x23c: mov ax,0xaa
   0x23f: push ax
   0x240: call _printf
   0x243: pop cx
5: exit(0);
   0x244: xor ax,ax
   0x246: push ax
   0x247: call _exit
   0x24a: pop cx
6: }
   0x24b: pop bp
   0x24c: ret
7: 0x24d:

```

Figure D.43: Line number program example: machine code

Suppose the line number program header includes the same values and resulting encoding illustrated in the previous Section D.5.2 on the preceding page.

Table D.2 on the next page shows one encoding of the line number program, which occupies 12 bytes.

## Appendix D. Examples (Informative)

Table D.2: Line number program example: one encoding

Opcode	Operand	Byte Stream
<a href="#">DW_LNS_advance_pc</a>	LEB128(0x239)	0x2, 0xb9, 0x04
SPECIAL† (2, 0)		0x12 (18 <sub>10</sub> )
SPECIAL† (2, 3)		0x36 (54 <sub>10</sub> )
SPECIAL† (1, 8)		0x71 (113 <sub>10</sub> )
SPECIAL† (1, 7)		0x65 (101 <sub>10</sub> )
<a href="#">DW_LNS_advance_pc</a>	LEB128(2)	0x2, 0x2
<a href="#">DW_LNE_end_sequence</a>		0x0, 0x1, 0x1

† The opcode notation SPECIAL(*m*,*n*) indicates the special opcode generated for a line advance of *m* and an operation advance of *n*.

1 Table D.3 shows an alternate encoding of the same program using standard  
 2 opcodes to advance the program counter; this encoding occupies 22 bytes.

Table D.3: Line number program example: alternate encoding

Opcode	Operand	Byte Stream
<a href="#">DW_LNS_fixed_advance_pc</a>	0x239	0x9, 0x39, 0x2
SPECIAL† (2, 0)		0x12 (18 <sub>10</sub> )
<a href="#">DW_LNS_fixed_advance_pc</a>	0x3	0x9, 0x3, 0x0
SPECIAL† (2, 0)		0x12 (18 <sub>10</sub> )
<a href="#">DW_LNS_fixed_advance_pc</a>	0x8	0x9, 0x8, 0x0
SPECIAL† (1, 0)		0x11 (17 <sub>10</sub> )
<a href="#">DW_LNS_fixed_advance_pc</a>	0x7	0x9, 0x7, 0x0
SPECIAL† (1, 0)		0x11 (17 <sub>10</sub> )
<a href="#">DW_LNS_fixed_advance_pc</a>	0x2	0x9, 0x2, 0x0
<a href="#">DW_LNE_end_sequence</a>		0x0, 0x1, 0x1

† SPECIAL is defined the same as in the preceding Table D.2.

## D.6 Call Frame Information Example

The following example uses a hypothetical RISC machine in the style of the Motorola 88000.

- Memory is byte addressed.
- Instructions are all 4 bytes each and word aligned.
- Instruction operands are typically of the form:
  - <destination.reg>, <source.reg>, <constant>
- The address for the load and store instructions is computed by adding the contents of the source register with the constant.
- There are eight 4-byte registers:

- R0 always 0
- R1 holds return address on call
- R2-R3 temp registers (not preserved on call)
- R4-R6 preserved on call
- R7 stack pointer

- The stack grows in the negative direction.
- The architectural ABI committee specifies that the stack pointer (R7) is the same as the CFA

Figure D.44 following shows two code fragments from a subroutine called foo that uses a frame pointer (in addition to the stack pointer). The first column values are byte addresses. <fs> denotes the stack frame size in bytes, namely 12.

An abstract table (see Section 6.4.1 on page 180) for the foo subroutine is shown in Table D.4 following. Corresponding fragments from the .debug\_frame section are shown in Table D.5 on page 345.

The following notations apply in Table D.4 on the next page:

1. R8 is the return address
2. s = same\_value rule
3. u = undefined rule
4. rN = register(N) rule
5. cN = offset(N) rule
6. a = architectural rule

## Appendix D. Examples (Informative)

```

        ;; start prologue
foo     sub   R7, R7, <fs>           ; Allocate frame
foo+4   store R1, R7, (<fs>-4)      ; Save the return address
foo+8   store R6, R7, (<fs>-8)      ; Save R6
foo+12  add   R6, R7, 0              ; R6 is now the Frame ptr
foo+16  store R4, R6, (<fs>-12)     ; Save a preserved reg
        ;; This subroutine does not change R5
        ...
        ;; Start epilogue (R7 is returned to entry value)
foo+64  load  R4, R6, (<fs>-12)     ; Restore R4
foo+68  load  R6, R7, (<fs>-8)      ; Restore R6
foo+72  load  R1, R7, (<fs>-4)      ; Restore return address
foo+76  add   R7, R7, <fs>          ; Deallocate frame
foo+80  jump  R1                    ; Return
foo+84

```

Figure D.44: Call frame information example: machine code fragments

Table D.4: Call frame information example: conceptual matrix

Location	CFA	R0	R1	R2	R3	R4	R5	R6	R7	R8
foo	[R7]+0	s	u	u	u	s	s	s	a	r1
foo+4	[R7]+fs	s	u	u	u	s	s	s	a	r1
foo+8	[R7]+fs	s	u	u	u	s	s	s	a	c-4
foo+12	[R7]+fs	s	u	u	u	s	s	c-8	a	c-4
foo+16	[R6]+fs	s	u	u	u	s	s	c-8	a	c-4
foo+20	[R6]+fs	s	u	u	u	c-12	s	c-8	a	c-4
...										
foo+64	[R6]+fs	s	u	u	u	c-12	s	c-8	a	c-4
foo+68	[R6]+fs	s	u	u	u	s	s	c-8	a	c-4
foo+72	[R7]+fs	s	u	u	u	s	s	s	a	c-4
foo+76	[R7]+fs	s	u	u	u	s	s	s	a	r1
foo+80	[R7]+0	s	u	u	u	s	s	s	a	r1

## Appendix D. Examples (Informative)

Table D.5: Call frame information example: common information entry encoding

Address	Value	Comment
cie	36	length
cie+4	0xffffffff	CIE_id
cie+8	4	version
cie+9	0	augmentation
cie+10	4	address size
cie+11	0	<i>Reserved</i>
cie+12	4	code_alignment_factor, <caf >
cie+13	-4	data_alignment_factor, <daf >
cie+14	8	R8 is the return addr.
cie+15	DW_CFA_def_cfa (7, 0)	CFA = [R7]+0
cie+18	DW_CFA_same_value (0)	R0 not modified (=0)
cie+20	DW_CFA_undefined (1)	R1 scratch
cie+22	DW_CFA_undefined (2)	R2 scratch
cie+24	DW_CFA_undefined (3)	R3 scratch
cie+26	DW_CFA_same_value (4)	R4 preserve
cie+28	DW_CFA_same_value (5)	R5 preserve
cie+30	DW_CFA_same_value (6)	R6 preserve
cie+32	DW_CFA_same_value (7)	R7 preserve
cie+34	DW_CFA_register (8, 1)	R8 is in R1
cie+37	DW_CFA_nop	padding
cie+38	DW_CFA_nop	padding
cie+39	DW_CFA_nop	padding
cie+40		

## Appendix D. Examples (Informative)

Table D.6: Call frame information example: frame description entry encoding

Address	Value	Comment†
fde	40	length
fde+4	cie	CIE_ptr
fde+8	foo	initial_location
fde+12	84	address_range
fde+16	DW_CFA_advance_loc(1)	instructions
fde+17	DW_CFA_def_cfa_offset(12)	<fs>
fde+19	DW_CFA_advance_loc(1)	4/<caf>
fde+20	DW_CFA_offset(8,1)	-4/<daf>(2nd parameter)
fde+22	DW_CFA_advance_loc(1)	
fde+23	DW_CFA_offset(6,2)	-8/<daf>(2nd parameter)
fde+25	DW_CFA_advance_loc(1)	
fde+26	DW_CFA_def_cfa_register(6)	
fde+28	DW_CFA_advance_loc(1)	
fde+29	DW_CFA_offset(4,3)	-12/<daf>(2nd parameter)
fde+31	DW_CFA_advance_loc(12)	44/<caf>
fde+32	DW_CFA_restore(4)	
fde+33	DW_CFA_advance_loc(1)	
fde+34	DW_CFA_restore(6)	
fde+35	DW_CFA_def_cfa_register(7)	
fde+37	DW_CFA_advance_loc(1)	
fde+38	DW_CFA_restore(8)	
fde+39	DW_CFA_advance_loc(1)	
fde+40	DW_CFA_def_cfa_offset(0)	
fde+42	DW_CFA_nop	padding
fde+43	DW_CFA_nop	padding
fde+44		

<sup>1</sup> †The following notations apply: <fs> = frame size, <caf> = code alignment factor, and <daf> = data alignment factor.

<sup>2</sup>



## 1 D.7 Inlining Examples

2 The pseudo-source in Figure D.45 following is used to illustrate the use of  
 3 DWARF to describe inlined subroutine calls. This example involves a nested  
 4 subprogram INNER that makes uplevel references to the formal parameter and  
 5 local variable of the containing subprogram OUTER.

```

inline procedure OUTER (OUTER_FORMAL : integer) =
  begin
    OUTER_LOCAL : integer;
    procedure INNER (INNER_FORMAL : integer) =
      begin
        INNER_LOCAL : integer;
        print(INNER_FORMAL + OUTER_LOCAL);
      end;
    INNER(OUTER_LOCAL);
    ...
    INNER(31);
  end;
! Call OUTER
!
OUTER(7);

```

Figure D.45: Inlining examples: pseudo-source fragment

6 There are several approaches that a compiler might take to inlining for this sort  
 7 of example. This presentation considers three such approaches, all of which  
 8 involve inline expansion of subprogram OUTER. (If OUTER is not inlined, the  
 9 inlining reduces to a simpler single level subset of the two level approaches  
 10 considered here.)

11 The approaches are:

- 12 1. Inline both OUTER and INNER in all cases
- 13 2. Inline OUTER, multiple INNERs  
 14 Treat INNER as a non-inlinable part of OUTER, compile and call a distinct  
 15 normal version of INNER defined within each inlining of OUTER.
- 16 3. Inline OUTER, one INNER  
 17 Compile INNER as a single normal subprogram which is called from every  
 18 inlining of OUTER.

19 This discussion does not consider why a compiler might choose one of these  
 20 approaches; it considers only how to describe the result.

## Appendix D. Examples (Informative)

1 In the examples that follow in this section, the debugging information entries are  
2 given mnemonic labels of the following form

3 `<io>.<ac>.<n>.<s>`

4 where

5 `<io>` is either INNER or OUTER to indicate to which subprogram the debugging  
6 information entry applies,

7 `<ac>` is either AI or CI to indicate “abstract instance” or “concrete instance”  
8 respectively,

9 `<n>` is the number of the alternative being considered, and

10 `<s>` is a sequence number that distinguishes the individual entries.

11 There is no implication that symbolic labels, nor any particular naming  
12 convention, are required in actual use.

13 For conciseness, declaration coordinates and call coordinates are omitted.

### 14 **D.7.1 Alternative #1: inline both OUTER and INNER**

15 A suitable abstract instance for an alternative where both OUTER and INNER are  
16 always inlined is shown in Figure [D.46 on the following page](#).

17 Notice in Figure [D.46](#) that the debugging information entry for INNER (labelled  
18 `INNER.AI.1.1$`) is nested in (is a child of) that for OUTER (labelled  
19 `OUTER.AI.1.1$`). Nonetheless, the abstract instance tree for INNER is considered  
20 to be separate and distinct from that for OUTER.

21 The call of OUTER shown in Figure [D.45 on the previous page](#) might be described  
22 as shown in Figure [D.47 on page 350](#).

### 23 **D.7.2 Alternative #2: Inline OUTER, multiple INNERs**

24 In the second alternative we assume that subprogram INNER is not inlinable for  
25 some reason, but subprogram OUTER is inlinable. Each concrete inlined instance  
26 of OUTER has its own normal instance of INNER. The abstract instance for OUTER,  
27 which includes INNER, is shown in Figure [D.48 on page 352](#).

28 Note that the debugging information in Figure [D.48](#) differs from that in Figure  
29 [D.46 on the following page](#) in that INNER lacks a `DW_AT_inline` attribute and  
30 therefore is not a distinct abstract instance. INNER is merely an out-of-line routine  
31 that is part of OUTER’s abstract instance. This is reflected in the Figure by the fact  
32 that the labels for INNER use the substring OUTER instead of INNER.

## Appendix D. Examples (Informative)

```
! Abstract instance for OUTER
!
OUTER.AI.1.1$:
  DW_TAG_subprogram
    DW_AT_name("OUTER")
    DW_AT_inline(DW_INL_declared_inlined)
    ! No low/high PCs
OUTER.AI.1.2$:
  DW_TAG_formal_parameter
    DW_AT_name("OUTER_FORMAL")
    DW_AT_type(reference to integer)
    ! No location
OUTER.AI.1.3$:
  DW_TAG_variable
    DW_AT_name("OUTER_LOCAL")
    DW_AT_type(reference to integer)
    ! No location
!
! Abstract instance for INNER
!
INNER.AI.1.1$:
  DW_TAG_subprogram
    DW_AT_name("INNER")
    DW_AT_inline(DW_INL_declared_inlined)
    ! No low/high PCs
INNER.AI.1.2$:
  DW_TAG_formal_parameter
    DW_AT_name("INNER_FORMAL")
    DW_AT_type(reference to integer)
    ! No location
INNER.AI.1.3$:
  DW_TAG_variable
    DW_AT_name("INNER_LOCAL")
    DW_AT_type(reference to integer)
    ! No location
...
0
! No DW_TAG_inlined_subroutine (concrete instance)
! for INNER corresponding to calls of INNER
...
0
```

Figure D.46: Inlining example #1: abstract instance

## Appendix D. Examples (Informative)

```
! Concrete instance for call "OUTER(7)"
!
OUTER.CI.1.1$:
  DW_TAG_inlined_subroutine
  ! No name
  DW_AT_abstract_origin(reference to OUTER.AI.1.1$)
  DW_AT_low_pc(...)
  DW_AT_high_pc(...)
OUTER.CI.1.2$:
  DW_TAG_formal_parameter
  ! No name
  DW_AT_abstract_origin(reference to OUTER.AI.1.2$)
  DW_AT_const_value(7)
OUTER.CI.1.3$:
  DW_TAG_variable
  ! No name
  DW_AT_abstract_origin(reference to OUTER.AI.1.3$)
  DW_AT_location(...)
!
! No DW_TAG_subprogram (abstract instance) for INNER
!
! Concrete instance for call INNER(OUTER_LOCAL)
!
INNER.CI.1.1$:
  DW_TAG_inlined_subroutine
  ! No name
  DW_AT_abstract_origin(reference to INNER.AI.1.1$)
  DW_AT_low_pc(...)
  DW_AT_high_pc(...)
  DW_AT_static_link(...)
INNER.CI.1.2$:
  DW_TAG_formal_parameter
  ! No name
  DW_AT_abstract_origin(reference to INNER.AI.1.2$)
  DW_AT_location(...)
INNER.CI.1.3$:
  DW_TAG_variable
  ! No name
  DW_AT_abstract_origin(reference to INNER.AI.1.3$)
  DW_AT_location(...)
...
0
! Another concrete instance of INNER within OUTER
! for the call "INNER(31)"
...
0
```

Figure D.47: Inlining example #1: concrete instance

## Appendix D. Examples (Informative)

1 A resulting concrete inlined instance of OUTER is shown in Figure D.49 on  
2 [page 354](#).

3 Notice in Figure D.49 that OUTER is expanded as a concrete inlined instance, and  
4 that INNER is nested within it as a concrete out-of-line subprogram. Because  
5 INNER is cloned for each inline expansion of OUTER, only the invariant attributes  
6 of INNER (for example, [DW\\_AT\\_name](#)) are specified in the abstract instance of  
7 OUTER, and the low-level, instance-specific attributes of INNER (for example,  
8 [DW\\_AT\\_low\\_pc](#)) are specified in each concrete instance of OUTER.

9 The several calls of INNER within OUTER are compiled as normal calls to the  
10 instance of INNER that is specific to the same instance of OUTER that contains the  
11 calls.

### 12 **D.7.3 Alternative #3: inline OUTER, one normal INNER**

13 In the third approach, one normal subprogram for INNER is compiled which is  
14 called from all concrete inlined instances of OUTER. The abstract instance for  
15 OUTER is shown in Figure [D.50 on page 355](#).

16 The most distinctive aspect of that Figure is that subprogram INNER exists only  
17 within the abstract instance of OUTER, and not in OUTER's concrete instance. In the  
18 abstract instance of OUTER, the description of INNER has the full complement of  
19 attributes that would be expected for a normal subprogram. While attributes  
20 such as [DW\\_AT\\_low\\_pc](#), [DW\\_AT\\_high\\_pc](#), [DW\\_AT\\_location](#), and so on,  
21 typically are omitted from an abstract instance because they are not invariant  
22 across instances of the containing abstract instance, in this case those same  
23 attributes are included precisely because they are invariant – there is only one  
24 subprogram INNER to be described and every description is the same.

25 A concrete inlined instance of OUTER is illustrated in Figure [D.51 on page 356](#).

26 Notice in Figure [D.51](#) that there is no DWARF representation for INNER at all; the  
27 representation of INNER does not vary across instances of OUTER and the abstract  
28 instance of OUTER includes the complete description of INNER, so that the  
29 description of INNER may be (and for reasons of space efficiency, should be)  
30 omitted from each concrete instance of OUTER.

## Appendix D. Examples (Informative)

```
! Abstract instance for OUTER
! abstract instance
OUTER.AI.2.1$:
  DW_TAG_subprogram
  DW_AT_name("OUTER")
  DW_AT_inline(DW_INL_declared_inlined)
  ! No low/high PCs
OUTER.AI.2.2$:
  DW_TAG_formal_parameter
  DW_AT_name("OUTER_FORMAL")
  DW_AT_type(reference to integer)
  ! No location
OUTER.AI.2.3$:
  DW_TAG_variable
  DW_AT_name("OUTER_LOCAL")
  DW_AT_type(reference to integer)
  ! No location
  !
  ! Nested out-of-line INNER subprogram
  !
OUTER.AI.2.4$:
  DW_TAG_subprogram
  DW_AT_name("INNER")
  ! No DW_AT_inline
  ! No low/high PCs, frame_base, etc.
OUTER.AI.2.5$:
  DW_TAG_formal_parameter
  DW_AT_name("INNER_FORMAL")
  DW_AT_type(reference to integer)
  ! No location
OUTER.AI.2.6$:
  DW_TAG_variable
  DW_AT_name("INNER_LOCAL")
  DW_AT_type(reference to integer)
  ! No location
  ...
  0
  ...
  0
```

Figure D.48: Inlining example #2: abstract instance

1 There is one aspect of this approach that is problematical from the DWARF  
2 perspective. The single compiled instance of INNER is assumed to access up-level  
3 variables of OUTER; however, those variables may well occur at varying positions  
4 within the frames that contain the concrete inlined instances. A compiler might  
5 implement this in several ways, including the use of additional

## Appendix D. Examples (Informative)

1 compiler-generated parameters that provide reference parameters for the  
2 up-level variables, or a compiler-generated static link like parameter that points  
3 to the group of up-level entities, among other possibilities. In either of these  
4 cases, the DWARF description for the location attribute of each uplevel variable  
5 needs to be different if accessed from within INNER compared to when accessed  
6 from within the instances of OUTER. An implementation is likely to require  
7 vendor-specific DWARF attributes and/or debugging information entries to  
8 describe such cases.

9 Note that in C++, a member function of a class defined within a function  
10 definition does not require any vendor-specific extensions because the C++  
11 language disallows access to entities that would give rise to this problem.  
12 (Neither extern variables nor static members require any form of static link for  
13 accessing purposes.)

## Appendix D. Examples (Informative)

```
! Concrete instance for call "OUTER(7)"
!  
OUTER.CI.2.1$:  
  DW_TAG_inlined_subroutine  
  ! No name  
  DW_AT_abstract_origin(reference to OUTER.AI.2.1$)  
  DW_AT_low_pc(...)  
  DW_AT_high_pc(...)  
OUTER.CI.2.2$:  
  DW_TAG_formal_parameter  
  ! No name  
  DW_AT_abstract_origin(reference to OUTER.AI.2.2$)  
  DW_AT_location(...)  
OUTER.CI.2.3$:  
  DW_TAG_variable  
  ! No name  
  DW_AT_abstract_origin(reference to OUTER.AI.2.3$)  
  DW_AT_location(...)  
!  
! Nested out-of-line INNER subprogram  
!  
OUTER.CI.2.4$:  
  DW_TAG_subprogram  
  ! No name  
  DW_AT_abstract_origin(reference to OUTER.AI.2.4$)  
  DW_AT_low_pc(...)  
  DW_AT_high_pc(...)  
  DW_AT_frame_base(...)  
  DW_AT_static_link(...)  
OUTER.CI.2.5$:  
  DW_TAG_formal_parameter  
  ! No name  
  DW_AT_abstract_origin(reference to OUTER.AI.2.5$)  
  DW_AT_location(...)  
OUTER.CI.2.6$:  
  DW_TAG_variable  
  ! No name  
  DW_AT_abstract_origin(reference to OUTER.AI.2.6$)  
  DW_AT_location(...)  
  ...  
  0  
  ...  
  0
```

Figure D.49: Inlining example #2: concrete instance



## Appendix D. Examples (Informative)

```
! Abstract instance for OUTER
!
OUTER.AI.3.1$:
  DW_TAG_subprogram
    DW_AT_name("OUTER")
    DW_AT_inline(DW_INL_declared_inlined)
    ! No low/high PCs
OUTER.AI.3.2$:
  DW_TAG_formal_parameter
    DW_AT_name("OUTER_FORMAL")
    DW_AT_type(reference to integer)
    ! No location
OUTER.AI.3.3$:
  DW_TAG_variable
    DW_AT_name("OUTER_LOCAL")
    DW_AT_type(reference to integer)
    ! No location
  !
  ! Normal INNER
  !
OUTER.AI.3.4$:
  DW_TAG_subprogram
    DW_AT_name("INNER")
    DW_AT_low_pc(...)
    DW_AT_high_pc(...)
    DW_AT_frame_base(...)
    DW_AT_static_link(...)
OUTER.AI.3.5$:
  DW_TAG_formal_parameter
    DW_AT_name("INNER_FORMAL")
    DW_AT_type(reference to integer)
    DW_AT_location(...)
OUTER.AI.3.6$:
  DW_TAG_variable
    DW_AT_name("INNER_LOCAL")
    DW_AT_type(reference to integer)
    DW_AT_location(...)
  ...
  0
  ...
  0
```

Figure D.50: Inlining example #3: abstract instance

## Appendix D. Examples (Informative)

```
! Concrete instance for call "OUTER(7)"
!  
OUTER.CI.3.1$:  
  DW_TAG_inlined_subroutine  
    ! No name  
    DW_AT_abstract_origin(reference to OUTER.AI.3.1$)  
    DW_AT_low_pc(...)  
    DW_AT_high_pc(...)  
    DW_AT_frame_base(...)  
OUTER.CI.3.2$:  
  DW_TAG_formal_parameter  
    ! No name  
    DW_AT_abstract_origin(reference to OUTER.AI.3.2$)  
    ! No type  
    DW_AT_location(...)  
OUTER.CI.3.3$:  
  DW_TAG_variable  
    ! No name  
    DW_AT_abstract_origin(reference to OUTER.AI.3.3$)  
    ! No type  
    DW_AT_location(...)  
    ! No DW_TAG_subprogram for "INNER"  
  ...  
0
```

Figure D.51: Inlining example #3: concrete instance

## 1 D.8 Constant Expression Example

2 C++ generalizes the notion of constant expressions to include constant  
3 expression user-defined literals and functions. The constant declarations in  
4 Figure D.52 can be represented as illustrated in Figure D.53 on the next page.

```
constexpr double mass = 9.8;  
constexpr int square (int x) { return x * x; }  
float arr[square(9)]; // square() called and inlined
```

Figure D.52: Constant expressions: C++ source

## Appendix D. Examples (Informative)

```
! For variable mass
!
1$: DW_TAG_const_type
    DW_AT_type(reference to "double")
2$: DW_TAG_variable
    DW_AT_name("mass")
    DW_AT_type(reference to 1$)
    DW_AT_const_expr(true)
    DW_AT_const_value(9.8)
! Abstract instance for square
!
10$: DW_TAG_subprogram
    DW_AT_name("square")
    DW_AT_type(reference to "int")
    DW_AT_inline(DW_INL_inlined)
11$: DW_TAG_formal_parameter
    DW_AT_name("x")
    DW_AT_type(reference to "int")
! Concrete instance for square(9)
!
20$: DW_TAG_inlined_subroutine
    DW_AT_abstract_origin(reference to 10$)
    DW_AT_const_expr(present)
    DW_AT_const_value(81)
    DW_TAG_formal_parameter
        DW_AT_abstract_origin(reference to 11$)
        DW_AT_const_value(9)
! Anonymous array type for arr
!
30$: DW_TAG_array_type
    DW_AT_type(reference to "float")
    DW_AT_byte_size(324) ! 81*4
    DW_TAG_subrange_type
        DW_AT_type(reference to "int")
        DW_AT_upper_bound(reference to 20$)
! Variable arr
!
40$: DW_TAG_variable
    DW_AT_name("arr")
    DW_AT_type(reference to 30$)
```

Figure D.53: Constant expressions: DWARF description

## 1 **D.9 Unicode Character Example**

2 The Unicode character encodings in Figure D.54 can be described in DWARF as  
3 illustrated in Figure D.55.

```
// C++ source
//
char16_t chr_a = u'h';
char32_t chr_b = U'h';
```

Figure D.54: Unicode character example: source

```
! DWARF description
!
1$: DW_TAG_base_type
   DW_AT_name("char16_t")
   DW_AT_encoding(DW_ATE_UTF)
   DW_AT_byte_size(2)
2$: DW_TAG_base_type
   DW_AT_name("char32_t")
   DW_AT_encoding(DW_ATE_UTF)
   DW_AT_byte_size(4)
3$: DW_TAG_variable
   DW_AT_name("chr_a")
   DW_AT_type(reference to 1$)
4$: DW_TAG_variable
   DW_AT_name("chr_b")
   DW_AT_type(reference to 2$)
```

Figure D.55: Unicode character example: DWARF description

## 1 D.10 Type-Safe Enumeration Example

2 The C++ type-safe enumerations in Figure D.56 can be described in DWARF as  
 3 illustrated in Figure D.57.

```
// C++ source
//
enum class E { E1, E2=100 };
E e1;
```

Figure D.56: Type-safe enumeration example: source

```
! DWARF description
!
11$: DW_TAG_enumeration_type
    DW_AT_name("E")
    DW_AT_type(reference to "int")
    DW_AT_enum_class(present)
12$: DW_TAG_enumerator
    DW_AT_name("E1")
    DW_AT_const_value(0)
13$: DW_TAG_enumerator
    DW_AT_name("E2")
    DW_AT_const_value(100)
14$: DW_TAG_variable
    DW_AT_name("e1")
    DW_AT_type(reference to 11$)
```

Figure D.57: Type-safe enumeration example: DWARF description

## 1 D.11 Template Examples

2 The C++ template example in Figure D.58 can be described in DWARF as  
 3 illustrated in Figure D.59.

```
// C++ source
//
template<class T>
struct wrapper {
    T comp;
};
wrapper<int> obj;
```

Figure D.58: C++ template example #1: source

```
! DWARF description
!
11$: DW_TAG_structure_type
    DW_AT_name("wrapper")
12$: DW_TAG_template_type_parameter
    DW_AT_name("T")
    DW_AT_type(reference to "int")
13$: DW_TAG_member
    DW_AT_name("comp")
    DW_AT_type(reference to 12$)
14$: DW_TAG_variable
    DW_AT_name("obj")
    DW_AT_type(reference to 11$)
```

Figure D.59: C++ template example #1: DWARF description

4 The actual type of the component `comp` is `int`, but in the DWARF the type  
 5 references the `DW_TAG_template_type_parameter` for `T`, which in turn  
 6 references `int`. This implies that in the original template `comp` was of type `T` and  
 7 that was replaced with `int` in the instance.

## Appendix D. Examples (Informative)

1 There exist situations where it is not possible for the DWARF to imply anything  
2 about the nature of the original template. Consider the C++ template source in  
3 Figure D.60 and the DWARF that can describe it in Figure D.61.

```
// C++ source
//
template<class T>
struct wrapper {
    T comp;
};
template<class U>
void consume(wrapper<U> formal)
{
    ...
}
wrapper<int> obj;
consume(obj);
```

Figure D.60: C++ template example #2: source

```
! DWARF description
!
11$: DW_TAG_structure_type
    DW_AT_name("wrapper")
12$: DW_TAG_template_type_parameter
    DW_AT_name("T")
    DW_AT_type(reference to "int")
13$: DW_TAG_member
    DW_AT_name("comp")
    DW_AT_type(reference to 12$)
14$: DW_TAG_variable
    DW_AT_name("obj")
    DW_AT_type(reference to 11$)
21$: DW_TAG_subprogram
    DW_AT_name("consume")
22$: DW_TAG_template_type_parameter
    DW_AT_name("U")
    DW_AT_type(reference to "int")
23$: DW_TAG_formal_parameter
    DW_AT_name("formal")
    DW_AT_type(reference to 11$)
```

Figure D.61: C++ template example #2: DWARF description

## Appendix D. Examples (Informative)

1 In the [DW\\_TAG\\_subprogram](#) entry for the instance of `consume`, `U` is described as  
2 `int`. The type of formal is `wrapper<U>` in the source. DWARF only represents  
3 instantiations of templates; there is no entry which represents `wrapper<U>` which  
4 is neither a template parameter nor a template instantiation. The type of formal is  
5 described as `wrapper<int>`, the instantiation of `wrapper<U>`, in the [DW\\_AT\\_type](#)  
6 attribute at 23\$. There is no description of the relationship between template type  
7 parameter `T` at 12\$ and `U` at 22\$ which was used to instantiate `wrapper<U>`.

8 A consequence of this is that the DWARF information would not distinguish  
9 between the existing example and one where the formal parameter of `consume`  
10 were declared in the source to be `wrapper<int>`.

### 11 D.12 Template Alias Examples

12 The C++ template alias shown in Figure [D.62](#) can be described in DWARF as  
13 illustrated in Figure [D.63 on the next page](#).

```
// C++ source, template alias example 1
//
template<typename T, typename U>
struct Alpha {
    T tango;
    U uniform;
};
template<typename V> using Beta = Alpha<V,V>;
Beta<long> b;
```

Figure D.62: C++ template alias example #1: source



## Appendix D. Examples (Informative)

```
! DWARF representation for variable 'b'
!  
20$: DW_TAG_structure_type  
    DW_AT_name("Alpha")  
21$: DW_TAG_template_type_parameter  
    DW_AT_name("T")  
    DW_AT_type(reference to "long")  
22$: DW_TAG_template_type_parameter  
    DW_AT_name("U")  
    DW_AT_type(reference to "long")  
23$: DW_TAG_member  
    DW_AT_name("tango")  
    DW_AT_type(reference to 21$)  
24$: DW_TAG_member  
    DW_AT_name("uniform")  
    DW_AT_type(reference to 22$)  
25$: DW_TAG_template_alias  
    DW_AT_name("Beta")  
    DW_AT_type(reference to 20$)  
26$: DW_TAG_template_type_parameter  
    DW_AT_name("V")  
    DW_AT_type(reference to "long")  
27$: DW_TAG_variable  
    DW_AT_name("b")  
    DW_AT_type(reference to 25$)
```

Figure D.63: C++ template alias example #1: DWARF description

1 Similarly, the C++ template alias shown in Figure [D.64](#) can be described in  
2 DWARF as illustrated in Figure [D.65 on the next page](#).

```
// C++ source, template alias example 2  
//  
template<class TX> struct X { };  
template<class TY> struct Y { };  
template<class T> using Z = Y<T>;  
X<Y<int>> y;  
X<Z<int>> z;
```

Figure D.64: C++ template alias example #2: source

## Appendix D. Examples (Informative)

```
! DWARF representation for X<Y<int>>
!
30$: DW_TAG_structure_type
      DW_AT_name("Y")
31$: DW_TAG_template_type_parameter
      DW_AT_name("TY")
      DW_AT_type(reference to "int")
32$: DW_TAG_structure_type
      DW_AT_name("X")
33$: DW_TAG_template_type_parameter
      DW_AT_name("TX")
      DW_AT_type(reference to 30$)
!
! DWARF representation for X<Z<int>>
!
40$: DW_TAG_template_alias
      DW_AT_name("Z")
      DW_AT_type(reference to 30$)
41$: DW_TAG_template_type_parameter
      DW_AT_name("T")
      DW_AT_type(reference to "int")
42$: DW_TAG_structure_type
      DW_AT_name("X")
43$: DW_TAG_template_type_parameter
      DW_AT_name("TX")
      DW_AT_type(reference to 40$)
!
! Note that 32$ and 42$ are actually the same type
!
50$: DW_TAG_variable
      DW_AT_name("y")
      DW_AT_type(reference to $32)
51$: DW_TAG_variable
      DW_AT_name("z")
      DW_AT_type(reference to $42)
```

Figure D.65: C++ template alias example #2: DWARF description

## 1 D.13 Implicit Pointer Examples

2 If the compiler determines that the value of an object is constant (either  
 3 throughout the program, or within a specific range), the compiler may choose to  
 4 materialize that constant only when used, rather than store it in memory or in a  
 5 register. The [DW\\_OP\\_implicit\\_value](#) operation can be used to describe such a  
 6 value. Sometimes, the value may not be constant, but still can be easily  
 7 rematerialized when needed. A DWARF expression terminating in  
 8 [DW\\_OP\\_stack\\_value](#) can be used for this case. The compiler may also eliminate  
 9 a pointer value where the target of the pointer resides in memory, and the  
 10 [DW\\_OP\\_stack\\_value](#) operator may be used to rematerialize that pointer value.  
 11 In other cases, the compiler will eliminate a pointer to an object that itself needs  
 12 to be materialized. Since the location of such an object cannot be represented as a  
 13 memory address, a DWARF expression cannot give either the location or the  
 14 actual value or a pointer variable that would refer to that object. The  
 15 [DW\\_OP\\_implicit\\_pointer](#) operation can be used to describe the pointer, and the  
 16 debugging information entry to which its first operand refers describes the value  
 17 of the dereferenced object. A DWARF consumer will not be able to show the  
 18 location or the value of the pointer variable, but it will be able to show the value  
 19 of the dereferenced pointer.

20 Consider the C source shown in Figure [D.66](#). Assume that the function `foo` is not  
 21 inlined, that the argument `x` is passed in register 5, and that the function `foo` is  
 22 optimized by the compiler into just an increment of the volatile variable `v`. Given  
 23 these assumptions a possible DWARF description is shown in Figure [D.67 on the](#)  
 24 [next page](#).

```

struct S { short a; char b, c; };
volatile int v;
void foo (int x)
{
    struct S s = { x, x + 2, x + 3 };
    char *p = &s.b;
    s.a++;
    v++;
}
int main ()
{
    foo (v+1);
    return 0;
}

```

Figure D.66: C implicit pointer example #1: source

## Appendix D. Examples (Informative)

```
1$: DW_TAG_structure_type
    DW_AT_name("S")
    DW_AT_byte_size(4)
10$: DW_TAG_member
    DW_AT_name("a")
    DW_AT_type(reference to "short int")
    DW_AT_data_member_location(constant 0)
11$: DW_TAG_member
    DW_AT_name("b")
    DW_AT_type(reference to "char")
    DW_AT_data_member_location(constant 2)
12$: DW_TAG_member
    DW_AT_name("c")
    DW_AT_type(reference to "char")
    DW_AT_data_member_location(constant 3)
2$: DW_TAG_subprogram
    DW_AT_name("foo")
20$: DW_TAG_formal_parameter
    DW_AT_name("x")
    DW_AT_type(reference to "int")
    DW_AT_location(DW_OP_reg5)
21$: DW_TAG_variable
    DW_AT_name("s")
    DW_AT_type(reference to S at 1$)
    DW_AT_location(expression=
        DW_OP_breg5(1) DW_OP_stack_value DW_OP_piece(2)
        DW_OP_breg5(2) DW_OP_stack_value DW_OP_piece(1)
        DW_OP_breg5(3) DW_OP_stack_value DW_OP_piece(1))
22$: DW_TAG_variable
    DW_AT_name("p")
    DW_AT_type(reference to "char *")
    DW_AT_location(expression=
        DW_OP_implicit_pointer(reference to 21$, 2))
```

Figure D.67: C implicit pointer example #1: DWARF description

1 In Figure D.67, even though variables `s` and `p` are both optimized away  
2 completely, this DWARF description still allows a debugger to print the value of  
3 the variable `s`, namely (2, 3, 4). Similarly, because the variable `s` does not live  
4 in memory, there is nothing to print for the value of `p`, but the debugger should  
5 still be able to show that `p[0]` is 3, `p[1]` is 4, `p[-1]` is 0 and `p[-2]` is 2.

## Appendix D. Examples (Informative)

1 As a further example, consider the C source shown in Figure D.68. Make the  
2 following assumptions about how the code is compiled:

- 3 • The function `foo` is inlined into function `main`
- 4 • The body of the `main` function is optimized to just three blocks of  
5 instructions which each increment the volatile variable `v`, followed by a  
6 block of instructions to return 0 from the function
- 7 • Label `label0` is at the start of the `main` function, `label1` follows the first `v++`  
8 block, `label2` follows the second `v++` block and `label3` is at the end of the  
9 `main` function
- 10 • Variable `b` is optimized away completely, as it isn't used
- 11 • The string literal `"opq"` is optimized away as well

12 Given these assumptions a possible DWARF description is shown in Figure D.69  
13 on the following page.

```
static const char *b = "opq";
volatile int v;
static inline void foo (int *p)
{
    (*p)++;
    v++;
    p++;
    (*p)++;
    v++;
}

int main ()
{
label0:
    int a[2] = 1, 2 ;
    v++;
label1:
    foo (a);
label2:
    return a[0] + a[1] - 5;
label3:
}
```

Figure D.68: C implicit pointer example #2: source

## Appendix D. Examples (Informative)

```
1$: DW_TAG_variable
    DW_AT_name("b")
    DW_AT_type(reference to "const char *")
    DW_AT_location(expression=
        DW_OP_implicit_pointer(reference to 2$, 0))
2$: DW_TAG_dwarf_procedure
    DW_AT_location(expression=
        DW_OP_implicit_value(4, {'o', 'p', 'q', '\0'}))
3$: DW_TAG_subprogram
    DW_AT_name("foo")
    DW_AT_inline(DW_INL_declared_inlined)
30$: DW_TAG_formal_parameter
    DW_AT_name("p")
    DW_AT_type(reference to "int *")
4$: DW_TAG_subprogram
    DW_AT_name("main")
40$: DW_TAG_variable
    DW_AT_name("a")
    DW_AT_type(reference to "int[2]")
    DW_AT_location(location list 98$)
41$: DW_TAG_inlined_subroutine
    DW_AT_abstract_origin(reference to 3$)
42$: DW_TAG_formal_parameter
    DW_AT_abstract_origin(reference to 30$)
    DW_AT_location(location list 99$)

! .debug_loclists section
98$: DW_LLE_start_end[<label0 in main> .. <label1 in main>)
    DW_OP_lit1 DW_OP_stack_value DW_OP_piece(4)
    DW_OP_lit2 DW_OP_stack_value DW_OP_piece(4)
    DW_LLE_start_end[<label1 in main> .. <label2 in main>)
    DW_OP_lit2 DW_OP_stack_value DW_OP_piece(4)
    DW_OP_lit2 DW_OP_stack_value DW_OP_piece(4)
    DW_LLE_start_end[<label2 in main> .. <label3 in main>)
    DW_OP_lit2 DW_OP_stack_value DW_OP_piece(4)
    DW_OP_lit3 DW_OP_stack_value DW_OP_piece(4)
    DW_LLE_end_of_list
99$: DW_LLE_start_end[<label1 in main> .. <label2 in main>)
    DW_OP_implicit_pointer(reference to 40$, 0)
    DW_LLE_start_end[<label2 in main> .. <label3 in main>)
    DW_OP_implicit_pointer(reference to 40$, 4)
    DW_LLE_end_of_list
```

Figure D.69: C implicit pointer example #2: DWARF description

## 1 D.14 String Type Examples

2 Consider the Fortran 2003 string type example source in Figure D.70 following.  
 3 The DWARF representation in Figure D.71 on the next page is appropriate.

```

program character_kind
  use iso_fortran_env
  implicit none
  integer, parameter :: ascii =
    selected_char_kind ("ascii")
  integer, parameter :: ucs4 =
    selected_char_kind ('ISO_10646')
  character(kind=ascii, len=26) :: alphabet
  character(kind=ucs4, len=30) :: hello_world
  character (len=*), parameter :: all_digits="0123456789"

  alphabet = ascii_"abcdefghijklmnopqrstuvwxyz"
  hello_world = ucs4_'Hello World and Ni Hao -- ' &
                // char (int (z'4F60'), ucs4)      &
                // char (int (z'597D'), ucs4)

  write (*,*) alphabet
  write (*,*) all_digits

  open (output_unit, encoding='UTF-8')
  write (*,*) trim (hello_world)
end program character_kind

```

Figure D.70: String type example: source

## Appendix D. Examples (Informative)

```
1$: DW_TAG_base_type
    DW_AT_encoding (DW_ATE_ASCII)

2$: DW_TAG_base_type
    DW_AT_encoding (DW_ATE_UCS)
    DW_AT_byte_size (4)

3$: DW_TAG_string_type
    DW_AT_byte_size (10)

4$: DW_TAG_const_type
    DW_AT_type (reference to 3$)

5$: DW_TAG_string_type
    DW_AT_type (1$)
    DW_AT_string_length ( ... )
    DW_AT_string_length_byte_size ( ... )
    DW_AT_data_location ( ... )

6$: DW_TAG_string_type
    DW_AT_type (2$)
    DW_AT_string_length ( ... )
    DW_AT_string_length_byte_size ( ... )
    DW_AT_data_location ( ... )

7$: DW_TAG_variable
    DW_AT_name (alphabet)
    DW_AT_type (5$)
    DW_AT_location ( ... )

8$: DW_TAG_constant
    DW_AT_name (all_digits)
    DW_AT_type (4$)
    DW_AT_const_value ( ... )

9$: DW_TAG_variable
    DW_AT_name (hello_world)
    DW_AT_type (6$)
    DW_AT_location ( ... )
```

Figure D.71: String type example: DWARF representation



## 1 D.15 Call Site Examples

2 The following examples use a hypothetical machine which:

- 3 • Passes the first argument in register 0, the second in register 1, and the third
- 4 in register 2.
- 5 • Keeps the stack pointer in register 3.
- 6 • Has one call preserved register 4.
- 7 • Returns a function value in register 0.

### 8 D.15.1 Call Site Example #1 (C)

9 Consider the C source in Figure D.72 following.

```
extern void fn1 (long int, long int, long int);

long int
fn2 (long int a, long int b, long int c)
{
    long int q = 2 * a;
    fn1 (5, 6, 7);
    return 0;
}

long int
fn3 (long int x, long int (*fn4) (long int *))
{
    long int v, w, w2, z;
    w = (*fn4) (&w2);
    v = (*fn4) (&w2);
    z = fn2 (1, v + 1, w);
    {
        int v1 = v + 4;
        z += fn2 (w, v * 2, x);
    }
    return z;
}
```

Figure D.72: Call Site Example #1: Source

10 Possible generated code for this source is shown using a suggestive pseudo-  
 11 assembly notation in Figure D.73 on the following page.

## Appendix D. Examples (Informative)

```
fn2:
L1:
    %reg2 = 7    ! Load the 3rd argument to fn1
    %reg1 = 6    ! Load the 2nd argument to fn1
    %reg0 = 5    ! Load the 1st argument to fn1
L2:
    call fn1
    %reg0 = 0    ! Load the return value from the function
    return
L3:
fn3:
    ! Decrease stack pointer to reserve local stack frame
    %reg3 = %reg3 - 32
    [%reg3] = %reg4    ! Save the call preserved register to
                        ! stack
    [%reg3 + 8] = %reg0 ! Preserve the x argument value
    [%reg3 + 16] = %reg1 ! Preserve the fn4 argument value
    %reg0 = %reg3 + 24 ! Load address of w2 as argument
    call %reg1        ! Call fn4 (indirect call)
L6:
    %reg2 = [%reg3 + 16] ! Load the fn4 argument value
    [%reg3 + 16] = %reg0 ! Save the result of the first call (w)
    %reg0 = %reg3 + 24 ! Load address of w2 as argument
    call %reg2        ! Call fn4 (indirect call)
L7:
    %reg4 = %reg0        ! Save the result of the second call (v)
                        ! into register.
    %reg2 = [%reg3 + 16] ! Load 3rd argument to fn2 (w)
    %reg1 = %reg4 + 1    ! Compute 2nd argument to fn2 (v + 1)
    %reg0 = 1            ! Load 1st argument to fn2
    call fn2
L4:
    %reg2 = [%reg3 + 8] ! Load the 3rd argument to fn2 (x)
    [%reg3 + 8] = %reg0 ! Save the result of the 3rd call (z)
    %reg0 = [%reg3 + 16] ! Load the 1st argument to fn2 (w)
    %reg1 = %reg4 + %reg4 ! Compute the 2nd argument to fn2 (v * 2)
    call fn2
L5:
    %reg2 = [%reg3 + 8] ! Load the value of z from the stack
    %reg0 = %reg0 + %reg2 ! Add result from the 4th call to it
L8:
    %reg4 = [%reg3]    ! Restore original value of call preserved
                        ! register
    %reg3 = %reg3 + 32 ! Leave stack frame
    return
```

Figure D.73: Call Site Example #1: Code

## Appendix D. Examples (Informative)

1 The location list for variable `a` in function `fn2` might look like the following  
2 (where the notation “*Range* [`m` .. `n`]” specifies the range of addresses from `m`  
3 through but not including `n` over which the following location description  
4 applies):

```
! Before the assignment to register 0, the argument a is live in register 0
!  
Range [L1 .. L2)  
    DW_OP_reg0  
  
! Afterwards, it is not. The value can perhaps be looked up in the caller  
!  
Range [L2 .. L3)  
    DW_OP_entry_value (1, DW_OP_reg0)  
    DW_OP_stack_value  
End-of-list
```

5 Similarly, the variable `q` in `fn2` then might have this location list:

```
! Before the assignment to register 0, the value of q can be computed as  
! two times the contents of register 0  
!  
Range [L1 .. L2)  
    DW_OP_lit2  
    DW_OP_breg0 0  
    DW_OP_mul  
    DW_OP_stack_value  
  
! Afterwards. it is not. It can be computed from the original value of  
! the first parameter, multiplied by two  
!  
Range [L2 .. L3)  
    DW_OP_lit2  
    DW_OP_entry_value (1, DW_OP_reg0)  
    DW_OP_mul  
    DW_OP_stack_value  
End-of-list
```

6 Variables `b` and `c` each have a location list similar to that for variable `a`, except for  
7 a different label between the two ranges and they use `DW_OP_reg1` and  
8 `DW_OP_reg2`, respectively, instead of `DW_OP_reg0`.

9 The call sites for all the calls in function `fn3` are children of the  
10 `DW_TAG_subprogram` entry for `fn3` (or of its `DW_TAG_lexical_block` entry if

## Appendix D. Examples (Informative)

1 there is any for the whole function). This is shown in Figure D.74.

*part 1 of 2*

```
DW_TAG_call_site
  DW_AT_call_return_pc(L6) ! First indirect call to (*fn4) in fn3.
  ! The address of the call is preserved across the call in memory at
  ! stack pointer + 16 bytes.
  DW_AT_call_target(DW_OP_breg3 16 DW_OP_deref)
  DW_TAG_call_site_parameter
    DW_AT_location(DW_OP_reg0)
    ! Value of the first parameter is equal to stack pointer + 24 bytes.
    DW_AT_call_value(DW_OP_breg3 24)
DW_TAG_call_site
  DW_AT_call_return_pc(L7) ! Second indirect call to (*fn4) in fn3.
  ! The address of the call is not preserved across the call anywhere, but
  ! could be perhaps looked up in fn3's caller.
  DW_AT_call_target(DW_OP_entry_value (1, DW_OP_reg1))
  DW_TAG_call_site_parameter
    DW_AT_location(DW_OP_reg0)
    DW_AT_call_value(DW_OP_breg3 24)
DW_TAG_call_site
  DW_AT_call_return_pc(L4) ! 3rd call in fn3, direct call to fn2
  DW_AT_call_origin(reference to fn2 DW_TAG_subprogram)
  DW_TAG_call_site_parameter
    DW_AT_call_parameter(reference to formal parameter a in subprogram fn2)
    DW_AT_location(DW_OP_reg0)
    ! First parameter to fn2 is constant 1
    DW_AT_call_value(DW_OP_lit1)
  DW_TAG_call_site_parameter
    DW_AT_call_parameter(reference to formal parameter b in subprogram fn2)
    DW_AT_location(DW_OP_reg1)
    ! Second parameter to fn2 can be computed as the value of the call
    ! preserved register 4 in the fn3 function plus one
    DW_AT_call_value(DW_OP_breg4 1)
  DW_TAG_call_site_parameter
    DW_AT_call_parameter(reference to formal parameter c in subprogram fn2)
    DW_AT_location(DW_OP_reg2)
    ! Third parameter's value is preserved in memory at fn3's stack pointer
    ! plus 16 bytes
    DW_AT_call_value(DW_OP_breg3 16 DW_OP_deref)
```

Figure D.74: Call site example #1: DWARF encoding

## Appendix D. Examples (Informative)

part 2 of 2

```
DW_TAG_lexical_block
  DW_AT_low_pc(L4)
  DW_AT_high_pc(L8)
  DW_TAG_variable
    DW_AT_name("v1")
    DW_AT_type(reference to int)
    ! Value of the v1 variable can be computed as value of register 4 plus 4
    DW_AT_location(DW_OP_breg4 4 DW_OP_stack_value)
  DW_TAG_call_site
    DW_AT_call_return_pc(L5) ! 4th call in fn3, direct call to fn2
    DW_AT_call_target(reference to subprogram fn2)
    DW_TAG_call_site_parameter
      DW_AT_call_parameter(reference to formal parameter a in subprogram fn2)
      DW_AT_location(DW_OP_reg0)
      ! Value of the 1st argument is preserved in memory at fn3's stack
      !   pointer + 16 bytes.
      DW_AT_call_value(DW_OP_breg3 16 DW_OP_deref)
    DW_TAG_call_site_parameter
      DW_AT_call_parameter(reference to formal parameter b in subprogram fn2)
      DW_AT_location(DW_OP_reg1)
      ! Value of the 2nd argument can be computed using the preserved
      !   register 4 multiplied by 2
      DW_AT_call_value(DW_OP_lit2 DW_OP_reg4 0 DW_OP_mul)
    DW_TAG_call_site_parameter
      DW_AT_call_parameter(reference to formal parameter c in subprogram fn2)
      DW_AT_location(DW_OP_reg2)
      ! Value of the 3rd argument is not preserved, but could be perhaps
      !   computed from the value passed fn3's caller.
      DW_AT_call_value(DW_OP_entry_value (1, DW_OP_reg0))
```

Figure D.74 Call site example #1: DWARF encoding (*concluded*)

1 **D.15.2 Call Site Example #2 (Fortran)**

2 Consider the Fortran source in Figure D.75 which is used to illustrate how  
3 Fortran's "pass by reference" parameters can be handled.

```
subroutine fn4 (n)
  integer :: n, x
  x = n
  n = n / 2
  call fn6
end subroutine
subroutine fn5 (n)
  interface fn4
    subroutine fn4 (n)
      integer :: n
    end subroutine
  end interface fn4
  integer :: n, x
  call fn4 (n)
  x = 5
  call fn4 (x)
end subroutine fn5
```

Figure D.75: Call site example #2: source

## Appendix D. Examples (Informative)

1 Possible generated code for this source is shown using a suggestive pseudo-  
2 assembly notation in Figure D.76.

```
fn4:
    %reg2 = [%reg0]    ! Load value of n (passed by reference)
    %reg2 = %reg2 / 2  ! Divide by 2
    [%reg0] = %reg2    ! Update value of n
    call fn6           ! Call some other function
    return

fn5:
    %reg3 = %reg3 - 8  ! Decrease stack pointer to create stack frame
    call fn4           ! Call fn4 with the same argument by reference
                        ! as fn5 has been called with

L9:
    [%reg3] = 5        ! Pass value of 5 by reference to fn4
    %reg0 = %reg3      ! Put address of the value 5 on the stack
                        ! into 1st argument register
    call fn4

L10:
    %reg3 = %reg3 + 8  ! Leave stack frame
    return
```

Figure D.76: Call site example #2: code

3 The location description for variable x in function fn4 might be:

```
DW_OP_entry_value 4 DW_OP_breg0 0 DW_OP_deref_size 4
DW_OP_stack_value
```

4 The call sites in (just) function fn5 might be as shown in Figure D.77 on the  
5 following page.

## Appendix D. Examples (Informative)

```
DW_TAG_call_site
  DW_AT_call_return_pc(L9)                ! First call to fn4
  DW_AT_call_origin(reference to subprogram fn4)
  DW_TAG_call_site_parameter
    DW_AT_call_parameter(reference to formal parameter n in subprogram fn4)
    DW_AT_location(DW_OP_reg0)
    ! The value of register 0 at the time of the call can be perhaps
    ! looked up in fn5's caller
    DW_AT_call_value(DW_OP_entry_value (1, DW_OP_reg0))
    ! DW_AT_call_data_location(DW_OP_push_object_address) ! left out, implicit
    ! And the actual value of the parameter can be also perhaps looked up in
    ! fn5's caller
    DW_AT_call_data_value(
      DW_OP_entry_value (4, DW_OP_breg0 0 DW_OP_deref_size 4))

DW_TAG_call_site
  DW_AT_call_return_pc(L10)               ! Second call to fn4
  DW_AT_call_origin(reference to subprogram fn4)
  DW_TAG_call_site_parameter
    DW_AT_call_parameter(reference to formal parameter n in subprogram fn4)
    DW_AT_location(DW_OP_reg0)
    ! The value of register 0 at the time of the call is equal to the stack
    ! pointer value in fn5
    DW_AT_call_value(DW_OP_breg3 0)
    ! DW_AT_call_data_location(DW_OP_push_object_address) ! left out, implicit
    ! And the value passed by reference is constant 5
    DW_AT_call_data_value(DW_OP_lit5)
```

Figure D.77: Call site example #2: DWARF encoding



## 1 D.16 Macro Example

2 Consider the C source in Figure D.78 following which is used to illustrate the  
3 DWARF encoding of macro information (see Section 6.3 on page 172).

*File a.c*

```
#include "a.h"
#define FUNCTION_LIKE_MACRO(x) 4+x
#include "b.h"
```

*File a.h*

```
#define LONGER_MACRO 1
#define B 2
#include "b.h"
#define B 3
```

*File b.h*

```
#undef B
#define D 3
#define FUNCTION_LIKE_MACRO(x) 4+x
```

Figure D.78: Macro example: source

4 Two possible encodings are shown. The first, in Figure D.79 on the following  
5 [page](#), is perhaps the simplest possible encoding. It includes all macro  
6 information from the main source file (a.c) as well as its two included files (a.h  
7 and b.h) in a single macro unit. Further, all strings are included as immediate  
8 operands of the macro operators (that is, there is no string pooling). The size of  
9 the macro unit is 160 bytes.

10 The second encoding, in Figure D.80 on page 381, saves space in two ways:

- 11 1. Longer strings are pooled by storing them in the `.debug_str` section where  
12 they can be referenced more than once.
- 13 2. Macro information entries contained in included files are represented as  
14 separate macro units which are then imported for each `#include` directive.

15 The combined size of the three macro units and their referenced strings is 129  
16 bytes.

## Appendix D. Examples (Informative)

```
! *** Section .debug_macro contents
! Macro unit for "a.c"
0$h:   Version:           5
      Flags:             2
           offset_size_flag: 0      ! 4-byte offsets
           debug_line_offset_flag: 1 ! Line number offset present
           opcode_operands_table_flag: 0 ! No extensions
      Offset in .debug_line section: 0 ! Line number offset
0$m:   DW_MACRO_start_file, 0, 0    ! Implicit Line: 0, File: 0 "a.c"
      DW_MACRO_start_file, 1, 1    ! #include Line: 1, File: 1 "a.h"
      DW_MACRO_define, 1, "LONGER_MACRO 1"
           ! #define Line: 1, String: "LONGER_MACRO 1"
      DW_MACRO_define, 2, "B 2"     ! #define Line: 2, String: "B 2"
      DW_MACRO_start_file, 3, 2    ! #include Line: 3, File: 2 "b.h"
      DW_MACRO_undef, 1, "B"       ! #undef Line: 1, String: "b"
      DW_MACRO_define, 2, "D 3"    ! #define Line: 2, String: "D 3"
      DW_MACRO_define, 3, "FUNCTION_LIKE_MACRO(x) 4+x"
           ! #define Line: 3,
           !   String: "FUNCTION_LIKE_MACRO(x) 4+x"
      DW_MACRO_end_file            ! End "b.h" -> back to "a.h"
      DW_MACRO_define, 4, "B 3"    ! #define Line: 4, String: "B 3"
      DW_MACRO_end_file            ! End "a.h" -> back to "a.c"
      DW_MACRO_define, 2, "FUNCTION_LIKE_MACRO(x) 4+x"
           ! #define Line: 2,
           !   String: "FUNCTION_LIKE_MACRO(x) 4+x"
      DW_MACRO_start_file, 3, 2    ! #include Line: 3, File: 2 "b.h"
      DW_MACRO_undef, 1, "B"       ! #undef Line: 1, String: "b"
      DW_MACRO_define, 2, "D 3"    ! #define Line: 2, String: "D 3"
      DW_MACRO_define, 3, "FUNCTION_LIKE_MACRO(x) 4+x"
           ! #define Line: 3,
           !   String: "FUNCTION_LIKE_MACRO(x) 4+x"
      DW_MACRO_end_file            ! End "b.h" -> back to "a.c"
      DW_MACRO_end_file            ! End "a.c" -> back to ""
      0                            ! End macro unit
```

Figure D.79: Macro example: simple DWARF encoding

## Appendix D. Examples (Informative)

```

! *** Section .debug_macro contents
! Macro unit for "a.c"
0$h:   Version:           5
      Flags:             2
           offset_size_flag: 0           ! 4-byte offsets
           debug_line_offset_flag: 1      ! Line number offset present
           opcode_operands_table_flag: 0 ! No extensions
      Offset in .debug_line section: 0    ! Line number offset
0$m:   DW_MACRO_start_file, 0, 0        ! Implicit Line: 0, File: 0 "a.c"
      DW_MACRO_start_file, 1, 1        ! #include Line: 1, File: 1 "a.h"
      DW_MACRO_import, i$1h           ! Import unit at i$1h (lines 1-2)
      DW_MACRO_start_file, 3, 2        ! #include Line: 3, File: 2 "b.h"
      DW_MACRO_import, i$2h           ! Import unit i$2h (lines all)
      DW_MACRO_end_file               ! End "b.h" -> back to "a.h"
      DW_MACRO_define, 4, "B 3"        ! #define Line: 4, String: "B 3"
      DW_MACRO_end_file               ! End "a.h" -> back to "a.c"
      DW_MACRO_define, 2, s$1          ! #define Line: 3,
      ! String: "FUNCTION_LIKE_MACRO(x) 4+x"
      DW_MACRO_start_file, 3, 2        ! #include Line: 3, File: 2 "b.h"
      DW_MACRO_import, i$2h           ! Import unit i$2h (lines all)
      DW_MACRO_end_file               ! End "b.h" -> back to "a.c"
      DW_MACRO_end_file               ! End "a.c" -> back to ""
      0                                ! End macro unit

! Macro unit for "a.h" lines 1-2
i$1h:  Version:           5
      Flags:             0
           offset_size_flag: 0           ! 4-byte offsets
           debug_line_offset_flag: 0      ! No line number offset
           opcode_operands_table_flag: 0 ! No extensions
i$1m:  DW_MACRO_define_strp, 1, s$2    ! #define Line: 1, String: "LONGER_MACRO 1"
      DW_MACRO_define, 2, "B 2"        ! #define Line: 2, String: "B 2"
      0                                ! End macro unit

! Macro unit for "b.h"
i$2h:  Version:           5
      Flags:             0
           offset_size_flag: 0           ! 4-byte offsets
           debug_line_offset_flag: 0      ! No line number offset
           opcode_operands_table_flag: 0 ! No extensions
i$2m:  DW_MACRO_undef, 1, "B"         ! #undef Line: 1, String: "B"
      DW_MACRO_define, 2, "D 3"        ! #define Line: 2, String: "D 3"
      DW_MACRO_define_strp, 3, s$1     ! #define Line: 3,
      ! String: "FUNCTION_LIKE_MACRO(x) 4+x"
      0                                ! End macro unit

! *** Section .debug_str contents
s$1:   String: "FUNCTION_LIKE_MACRO(x) 4+x"
s$2:   String: "LONGER_MACRO 1"

```

Figure D.80: Macro example: sharable DWARF encoding

## Appendix D. Examples (Informative)

1 A number of observations are worth mentioning:

- 2 • Strings that are the same size as a reference or less are better represented as  
3 immediate operands. Strings longer than twice the size of a reference are  
4 better stored in the string table if there are at least two references.
- 5 • There is a trade-off between the size of the macro information of a file and  
6 the number of times it is included when evaluating whether to create a  
7 separate macro unit. However, the amount of overhead (the size of a macro  
8 header) needed to represent a unit as well as the size of the operation to  
9 import a macro unit are both small.
- 10 • A macro unit need not describe all of the macro information in a file. For  
11 example, in Figure D.80 the second macro unit (beginning at i\$1h) includes  
12 macros from just the first two lines of file a.h.
- 13 • An implementation may be able to share macro units across object files (not  
14 shown in this example). To support this, it may be advantageous to create  
15 macro units in cases where they do not offer an advantage in a single  
16 compilation of itself.
- 17 • The header of a macro unit that contains a [DW\\_MACRO\\_start\\_file](#)  
18 operation must include a reference to the compilation line number header  
19 to allow interpretation of the file number operands in those commands.  
20 However, the presence of those offsets complicates or may preclude sharing  
21 across compilations.

## D.17 Parameter Default Value Examples

The default expressions for parameters `x` and `y` in the C++ function declaration in Figure D.81 can be described in DWARF as illustrated in Figure D.82.

```
void g (int x = 13;
       int y = f());
```

Figure D.81: Default value example #1: C++ source

```
DW_TAG_subprogram
  DW_AT_name ("g")

  DW_TAG_formal_parameter
    DW_AT_name ("x")
    DW_AT_type (reference to type "int")
    DW_AT_default_value@DW_FORM_sdata (13)

  DW_TAG_formal_parameter
    DW_AT_name ("y")
    DW_AT_type (reference to type "int")
    DW_AT_default_value@DW_FORM_string ("f()")
```

Figure D.82: Default value example #1: DWARF encoding

In Figure D.82, note the following:

1. This figure explicitly shows the form used by certain attributes (indicated by a trailing `@DW_FORM_XXX`) when it is critical, while the form is left implicit in most other examples.
2. The string value for `y` is three characters in length and does not include any quotes. (The quotes are an artifact of the presumed dumper tool that created this interpretation.)
3. The default value for `x` could also be encoded as `DW_AT_default_value@DW_FORM_string("13")`; however, this is generally a less convenient form and less efficient for consumers to process. a less convenient form and less efficient for consumers to process.

## Appendix D. Examples (Informative)

1 A string form in `DW_AT_default_value` always represents a source code  
2 fragment, even in languages that have a native string type. For example, the  
3 default string parameter of the Ada function in Figure D.83 is encoded in  
4 DWARF as a string containing the Ada string literal, including the source  
5 quotation marks, as shown in Figure D.84.

```
procedure s (x : string := "abc";
            y : string := "abcd"+10) is
begin
end s;
```

Figure D.83: Default value example #2: Ada source

```
DW_TAG_subprogram
  DW_AT_name ("s")

  DW_TAG_formal_parameter
    DW_AT_name ("x")
    DW_AT_type (reference to type "string")
    DW_AT_default_value@DW_FORM_data4 (0x61626364)    ! Big-endian

  DW_TAG_formal_parameter
    DW_AT_name ("y")
    DW_AT_type (reference to type "string")
    DW_AT_default_value@DW_FORM_string ("abcd"+10)
```

Figure D.84: Default value example #2: DWARF encoding

Appendix D. Examples (Informative)

*(empty page)*

# Appendix E

## DWARF Compression and Duplicate Elimination (Informative)

DWARF can use a lot of disk space.

This is especially true for C++, where the depth and complexity of headers can mean that many, many (possibly thousands of) declarations are repeated in every compilation unit. C++ templates can also mean that some functions and their DWARF descriptions get duplicated.

This Appendix describes techniques for using the DWARF representation in combination with features and characteristics of some common object file representations to reduce redundancy without losing information. It is worth emphasizing that none of these techniques are necessary to provide a complete and accurate DWARF description; they are solely concerned with reducing the size of DWARF information.

The techniques described here depend more directly and more obviously on object file concepts and linker mechanisms than most other parts of DWARF. While the presentation tends to use the vocabulary of specific systems, this is primarily to aid in describing the techniques by appealing to well-known terminology. These techniques can be employed on any system that supports certain general functional capabilities (described below).

### E.1 Using Compilation Units

#### E.1.1 Overview

The general approach is to break up the debug information of a compilation into separate normal and partial compilation units, each consisting of one or more



## Appendix E. Compression (Informative)

1 sections. By arranging that a sufficiently similar partitioning occurs in other  
2 compilations, a suitable system linker can delete redundant groups of sections  
3 when combining object files.

4 *The following uses some traditional section naming here but aside from the DWARF*  
5 *sections, the names are just meant to suggest traditional contents as a way of explaining*  
6 *the approach, not to be limiting.*

7 A traditional relocatable object output file from a single compilation might  
8 contain sections named:

```
9     .data  
10    .text  
11    .debug_info  
12    .debug_abbrev  
13    .debug_line  
14    .debug_aranges
```

15 A relocatable object file from a compilation system attempting duplicate DWARF  
16 elimination might contain sections as in:

```
17    .data  
18    .text  
19    .debug_info  
20    .debug_abbrev  
21    .debug_line  
22    .debug_aranges
```

23 followed (or preceded, the order is not significant) by a series of section groups:

```
24    ==== Section group 1  
25        .debug_info  
26        .debug_abbrev  
27        .debug_line  
28    ==== ...  
29    ==== Section group N  
30        .debug_info  
31        .debug_abbrev  
32        .debug_line
```

33 where each section group might or might not contain executable code (.text  
34 sections) or data (.data sections).

## Appendix E. Compression (Informative)

1 A *section group* is a named set of section contributions within an object file with  
2 the property that the entire set of section contributions must be retained or  
3 discarded as a whole; no partial elimination is allowed. Section groups can  
4 generally be handled by a linker in two ways:

- 5 1. Given multiple identical (duplicate) section groups, one of them is chosen to  
6 be kept and used, while the rest are discarded.
- 7 2. Given a section group that is not referenced from any section outside of the  
8 section group, the section group is discarded.

9 Which handling applies may be indicated by the section group itself and/or  
10 selection of certain linker options.

11 For example, if a linker determines that section group 1 from A.o and section  
12 group 3 from B.o are identical, it could discard one group and arrange that all  
13 references in A.o and B.o apply to the remaining one of the two identical section  
14 groups. This saves space.

15 An important part of making it possible to “redirect” references to the surviving  
16 section group is the use of consistently chosen linker global symbols for referring  
17 to locations within each section group. It follows that references are simply to  
18 external names and the linker already knows how to match up references and  
19 definitions.

20 What is minimally needed from the object file format and system linker (outside  
21 of DWARF itself, and normal object/linker facilities such as simple relocations)  
22 are:

- 23 1. A means to reference the `.debug_info` information of one compilation unit  
24 from the `.debug_info` section of another compilation unit  
25 (`DW_FORM_ref_addr` provides this).
- 26 2. A means to combine multiple contributions to specific sections (for example,  
27 `.debug_info`) into a single object file.
- 28 3. A means to identify a section group (giving it a name).
- 29 4. A means to indicate which sections go together to make up a section group,  
30 so that the group can be treated as a unit (kept or discarded).
- 31 5. A means to indicate how each section group should be processed by the  
32 linker.

33 *The notion of section and section contribution used here corresponds closely to the*  
34 *similarly named concepts in the ELF object file representation. The notion of section*  
35 *group is an abstraction of common extensions of the ELF representation widely known as*

1 “COMDATs” or “COMDAT sections.” (Other object file representations provide  
2 COMDAT-style mechanisms as well.) There are several variations in the COMDAT  
3 schemes in common use, any of which should be sufficient for the purposes of the  
4 DWARF duplicate elimination techniques described here.

## 5 E.1.2 Naming and Usage Considerations

6 A precise description of the means of deriving names usable by the linker to  
7 access DWARF entities is not part of this specification. Nonetheless, an outline of  
8 a usable approach is given here to make this more understandable and to guide  
9 implementors.

10 Implementations should clearly document their naming conventions.

11 In the following, it will be helpful to refer to the examples in Figure E.1 through  
12 Figure E.8 of Section E.1.3 on page 392.

### 13 Section Group Names

14 Section groups must have a section group name. For the subsequent C++  
15 example, a name like

16 `<producer-prefix>.<file-designator>.<gid-number>`

17 will suffice, where

18 `<producer-prefix>` is some string specific to the producer, which has a  
19 language-designation embedded in the name when appropriate.  
20 (Alternatively, the language name could be embedded in the  
21 `<gid-number>`).

22 `<file-designator>` names the file, such as `wa.h` in the example.

23 `<gid-number>` is a string generated to identify the specific `wa.h` header file in  
24 such a way that

- 25 • a ‘matching’ output from another compile generates the same  
26 `<gid-number>`, and
- 27 • a non-matching output (say because of `#defines`) generates a different  
28 `<gid-number>`.

29 *It may be useful to think of a `<gid-number>` as a kind of “digital signature” that allows a*  
30 *fast test for the equality of two section groups.*

31 So, for example, the section group corresponding to file `wa.h` above is given the  
32 name `my.compiler.company.cpp.wa.h.123456`.

## 1 **Debugging Information Entry Names**

2 Global labels for debugging information entries (the need for which is explained  
3 below) within a section group can be given names of the form

4 `<prefix>.<file-designator>.<gid-number>.<die-number>`

5 such as

6 `my.compiler.company.wa.h.123456.987`

7 where

8 `<prefix>` distinguishes this as a DWARF debug info name, and should identify  
9 the producer and, when appropriate, the language.

10 `<file-designator>` and `<gid-number>` are as above.

11 `<die-number>` could be a number sequentially assigned to entities (tokens,  
12 perhaps) found during compilation.

13 In general, every point in the section group `.debug_info` that could be referenced  
14 from outside by *any* compilation unit must normally have an external name  
15 generated for it in the linker symbol table, whether the current compilation  
16 references all those points or not.

17 *The completeness of the set of names generated is a quality-of-implementation issue.*

18 It is up to the producer to ensure that if `<die-numbers>` in separate compilations  
19 would not match properly then a distinct `<gid-number>` is generated.

20 Note that only section groups that are designated as `duplicate-removal-applies`  
21 actually require the

22 `<prefix>.<file-designator>.<gid-number>.<die-number>`

23 external labels for debugging information entries as all other section group  
24 sections can use 'local' labels (section-relative relocations).

25 (This is a consequence of separate compilation, not a rule imposed by this  
26 document.)

27 *Local labels use references with form `DW_FORM_ref4` or `DW_FORM_ref8`. (These are  
28 affected by relocations so `DW_FORM_ref_udata`, `DW_FORM_ref1` and  
29 `DW_FORM_ref2` are normally not usable and `DW_FORM_ref_addr` is not necessary for  
30 a local label.)*

1 **E.1.2.1 Use of DW\_TAG\_compile\_unit versus DW\_TAG\_partial\_unit**

2 A section group compilation unit that uses [DW\\_TAG\\_compile\\_unit](#) is like any  
3 other compilation unit, in that its contents are evaluated by consumers as though  
4 it were an ordinary compilation unit.

5 An #include directive appearing outside any other declarations is a good  
6 candidate to be represented using [DW\\_TAG\\_compile\\_unit](#). However, an  
7 #include appearing inside a C++ namespace declaration or a function, for  
8 example, is not a good candidate because the entities included are not necessarily  
9 file level entities.

10 This also applies to Fortran INCLUDE lines when declarations are included into  
11 a subprogram or module context.

12 Consequently a compiler must use [DW\\_TAG\\_partial\\_unit](#) (instead of  
13 [DW\\_TAG\\_compile\\_unit](#)) in a section group whenever the section group contents  
14 are not necessarily globally visible. This directs consumers to ignore that  
15 compilation unit when scanning top level declarations and definitions.

16 The [DW\\_TAG\\_partial\\_unit](#) compilation unit will be referenced from elsewhere  
17 and the referencing locations give the appropriate context for interpreting the  
18 partial compilation unit.

19 A [DW\\_TAG\\_partial\\_unit](#) entry may have, as appropriate, any of the attributes  
20 assigned to a [DW\\_TAG\\_compile\\_unit](#).

21 **E.1.2.2 Use of DW\_TAG\_imported\_unit**

22 A [DW\\_TAG\\_imported\\_unit](#) debugging information entry has an  
23 [DW\\_AT\\_import](#) attribute referencing a [DW\\_TAG\\_compile\\_unit](#) or  
24 [DW\\_TAG\\_partial\\_unit](#) debugging information entry.

25 A [DW\\_TAG\\_imported\\_unit](#) debugging information entry refers to a  
26 [DW\\_TAG\\_compile\\_unit](#) or [DW\\_TAG\\_partial\\_unit](#) debugging information entry  
27 to specify that the [DW\\_TAG\\_compile\\_unit](#) or [DW\\_TAG\\_partial\\_unit](#) contents  
28 logically appear at the point of the [DW\\_TAG\\_imported\\_unit](#) entry.

29 **E.1.2.3 Use of DW\_FORM\_ref\_addr**

30 Use [DW\\_FORM\\_ref\\_addr](#) to reference from one compilation unit's debugging  
31 information entries to those of another compilation unit.

## Appendix E. Compression (Informative)

1 When referencing into a removable section group `.debug_info` from another  
2 `.debug_info` (from anywhere), the

3 `<prefix>.<file-designator>.<gid-number>.<die-number>`

4 name should be used for an external symbol and a relocation generated based on  
5 that name.

6 *When referencing into a non-section group `.debug_info`, from another `.debug_info`  
7 (from anywhere) `DW_FORM_ref_addr` is still the form to be used, but a section-relative  
8 relocation generated by use of a non-exported name (often called an “internal name”)  
9 may be used for references within the same object file.*

### 10 E.1.3 Examples

11 This section provides several examples in order to have a concrete basis for  
12 discussion.

13 In these examples, the focus is on the arrangement of DWARF information into  
14 sections (specifically the `.debug_info` section) and the naming conventions used  
15 to achieve references into section groups. In practice, all of the examples that  
16 follow involve DWARF sections other than just `.debug_info` (for example,  
17 `.debug_line`, `.debug_aranges`, or others); however, only the `.debug_info` section  
18 is shown to keep the examples compact and easier to read.

19 The grouping of sections into a named set is shown, but the means for achieving  
20 this in terms of the underlying object language is not (and varies from system to  
21 system).

#### 22 E.1.3.1 C++ Example

23 The C++ source in Figure [E.1 on the following page](#) is used to illustrate the  
24 DWARF representation intended to allow duplicate elimination.

25 Figure [E.2 on the next page](#) shows the section group corresponding to the  
26 included file `wa.h`.

27 Figure [E.3 on page 394](#) shows the “normal” DWARF sections, which are not part  
28 of any section group, and how they make use of the information in the section  
29 group shown above.

30 This example uses `DW_TAG_compile_unit` for the section group, implying that  
31 the contents of the compilation unit are globally visible (in accordance with C++  
32 language rules). `DW_TAG_partial_unit` is not needed for the same reason.

## Appendix E. Compression (Informative)

*File wa.h*

```
struct A {
    int i;
};
```

*File wa.c*

```
#include "wa.h";
int
f(A &a)
{
    return a.i + 2;
}
```

Figure E.1: Duplicate elimination example #1: C++ Source

```
==== Section group name:
    my.compiler.company.cpp.wa.h.123456
== section .debug_info
DW.cpp.wa.h.123456.1:      ! linker global symbol
    DW_TAG_compile_unit
        DW_AT_language_name(DW_LNAME_C_plus_plus)
        ... ! other unit attributes
DW.cpp.wa.h.123456.2:      ! linker global symbol
    DW_TAG_base_type
        DW_AT_name("int")
DW.cpp.wa.h.123456.3:      ! linker global symbol
    DW_TAG_structure_type
        DW_AT_name("A")
DW.cpp.wa.h.123456.4:      ! linker global symbol
    DW_TAG_member
        DW_AT_name("i")
        DW_AT_type(DW_FORM_ref<n> to DW.cpp.wa.h.123456.2)
        ! (This is a local reference, so the more
        ! compact form DW_FORM_ref<n>
        ! for n = 1,2,4, or 8 can be used)
```

Figure E.2: Duplicate elimination example #1: DWARF section group

### 1 E.1.3.2 C Example

2 The C++ example in this Section might appear to be equally valid as a C  
3 example. However, for C it is prudent to include a `DW_TAG_imported_unit` in  
4 the primary unit (see Figure E.3 on the next page) as well as an `DW_AT_import`  
5 attribute that refers to the proper unit in the section group.

## Appendix E. Compression (Informative)

```
== section .text
    [generated code for function f]
== section .debug_info
    DW_TAG_compile_unit
.L1:                                ! local (non-linker) symbol
    DW_TAG_reference_type
        DW_AT_type(reference to DW.cpp.wa.h.123456.3)
    DW_TAG_subprogram
        DW_AT_name("f")
        DW_AT_type(reference to DW.cpp.wa.h.123456.2)
    DW_TAG_variable
        DW_AT_name("a")
        DW_AT_type(reference to .L1)
    ...
```

Figure E.3: Duplicate elimination example #1: primary compilation unit

1 *The C rules for consistency of global (file scope) symbols across compilations are less*  
2 *strict than for C++; inclusion of the import unit attribute assures that the declarations of*  
3 *the proper section group are considered before declarations from other compilations.*

### 4 **E.1.3.3 Fortran Example**

5 For a Fortran example, consider Figure E.4.

*File CommonStuff.fh*

```
IMPLICIT INTEGER(A-Z)
COMMON /Common1/ C(100)
PARAMETER(SEVEN = 7)
```

*File Func.f*

```
FUNCTION FOO (N)
INCLUDE 'CommonStuff.fh'
FOO = C(N + SEVEN)
RETURN
END
```

Figure E.4: Duplicate elimination example #2: Fortran source

6 Figure E.5 on the next page shows the section group corresponding to the  
7 included file CommonStuff.fh.

8 Figure E.6 on page 396 shows the sections for the primary compilation unit.

9 A companion main program is shown in Figure E.7 on page 396



## Appendix E. Compression (Informative)

```
==== Section group name:

    my.f90.company.f90.CommonStuff.fh.654321

== section .debug_info

DW.myf90.CommonStuff.fh.654321.1:    ! linker global symbol
    DW_TAG_partial_unit
        ! ...compilation unit attributes, including...
        DW_AT_language_name(DW_LNAME_Fortran)
        DW_AT_identifier_case(DW_ID_case_insensitive)

DW.myf90.CommonStuff.fh.654321.2:    ! linker global symbol
3$: DW_TAG_array_type
    ! unnamed
    DW_AT_type(reference to DW.f90.F90$main.f.2)
        ! base type INTEGER
    DW_TAG_subrange_type
        DW_AT_type(reference to DW.f90.F90$main.f.2)
            ! base type INTEGER
        DW_AT_lower_bound(constant 1)
        DW_AT_upper_bound(constant 100)

DW.myf90.CommonStuff.fh.654321.3:    ! linker global symbol
    DW_TAG_common_block
        DW_AT_name("Common1")
        DW_AT_location(Address of common block Common1)
    DW_TAG_variable
        DW_AT_name("C")
        DW_AT_type(reference to 3$)
        DW_AT_location(address of C)

DW.myf90.CommonStuff.fh.654321.4:    ! linker global symbol
    DW_TAG_constant
        DW_AT_name("SEVEN")
        DW_AT_type(reference to DW.f90.F90$main.f.2)
            ! base type INTEGER
        DW_AT_const_value(constant 7)
```

Figure E.5: Duplicate elimination example #2: DWARF section group

## Appendix E. Compression (Informative)

```
== section .text
    [code for function Foo]

== section .debug_info
    DW_TAG_compile_unit
        DW_TAG_subprogram
            DW_AT_name("Foo")
            DW_AT_type(reference to DW.f90.F90$main.f.2)
                ! base type INTEGER
            DW_TAG_imported_unit
                DW_AT_import(reference to
                    DW.myf90.CommonStuff.fh.654321.1)
            DW_TAG_common_inclusion ! For Common1
                DW_AT_common_reference(reference to
                    DW.myf90.CommonStuff.fh.654321.3)
            DW_TAG_variable ! For function result
                DW_AT_name("Foo")
                DW_AT_type(reference to DW.f90.F90$main.f.2)
                    ! base type INTEGER
```

Figure E.6: Duplicate elimination example #2: primary unit

*File Main.f*

```
INCLUDE 'CommonStuff.fh'
C(50) = 8
PRINT *, 'Result = ', F00(50 - SEVEN)
END
```

Figure E.7: Duplicate elimination example #2: companion source

1 That main program results in an object file that contained a duplicate of the  
2 section group named `my.f90.company.f90.CommonStuff.fh.654321`  
3 corresponding to the included file as well as the remainder of the main  
4 subprogram as shown in [Figure E.8 on the following page](#).

5 This example uses `DW_TAG_partial_unit` for the section group because the  
6 included declarations are not independently visible as global entities.

## 7 E.2 Using Type Units

8 A large portion of debug information is type information, and in a typical  
9 compilation environment, many types are duplicated many times. One method  
10 of controlling the amount of duplication is separating each type into a separate  
11 COMDAT `.debug_info` section and arranging for the linker to recognize and

## Appendix E. Compression (Informative)

```
== section .debug_info
  DW_TAG_compile_unit
    DW_AT_name(F90$main)
    DW_TAG_base_type
      DW_AT_name("INTEGER")
      DW_AT_encoding(DW_ATE_signed)
      DW_AT_byte_size(...)

    DW_TAG_base_type
      ...
    ... ! other base types
  DW_TAG_subprogram
    DW_AT_name("F90$main")
    DW_TAG_imported_unit
      DW_AT_import(reference to
        DW.myf90.CommonStuff.fh.654321.1)
    DW_TAG_common_inclusion ! for Common1
      DW_AT_common_reference(reference to
        DW.myf90.CommonStuff.fh.654321.3)
    ...
```

Figure E.8: Duplicate elimination example #2: companion DWARF

1 eliminate duplicates at the individual type level.

2 Using this technique, each substantial type definition is placed in its own  
3 individual section, while the remainder of the DWARF information (non-type  
4 information, incomplete type declarations, and definitions of trivial types) is  
5 placed in the usual debug information section. In a typical implementation, the  
6 relocatable object file may contain one of each of these debug sections:

7 .debug\_abbrev  
8 .debug\_info  
9 .debug\_line

10 and any number of additional COMDAT .debug\_info sections containing type  
11 units.

## Appendix E. Compression (Informative)

1 As discussed in the previous section (Section [E.1 on page 386](#)), many linkers  
2 today support the concept of a COMDAT group or linkonce section. The general  
3 idea is that a “key” can be attached to a section or a group of sections, and the  
4 linker will include only one copy of a section group (or individual section) for  
5 any given key. For COMDAT `.debug_info` sections, the key is the type signature  
6 formed from the algorithm given in Section [7.32 on page 257](#).

### 7 **E.2.1 Signature Computation Example**

8 As an example, consider a C++ header file containing the type definitions shown  
9 in Figure [E.9](#).

```
namespace N {  
  
    struct B;  
  
    struct C {  
        int x;  
        int y;  
    };  
  
    class A {  
    public:  
        A(int v);  
        int v();  
    private:  
        int v_;  
        struct A *next;  
        struct B *bp;  
        struct C c;  
    };  
}
```

Figure E.9: Type signature examples: C++ source

10 Next, consider one possible representation of the DWARF information that  
11 describes the type “struct C” as shown in [E.10 on the following page](#).

12 In computing a signature for the type `N::C`, flatten the type description into a  
13 byte stream according to the procedure outlined in Section [7.32 on page 257](#). The  
14 result is shown in Figure [E.11 on page 400](#).

## Appendix E. Compression (Informative)

```
DW_TAG_type_unit
  DW_AT_language_name : DW_LNAME_C_plus_plus (4)
  DW_TAG_namespace
    DW_AT_name : "N"
L1:
  DW_TAG_structure_type
    DW_AT_name : "C"
    DW_AT_byte_size : 8
    DW_AT_decl_file : 1
    DW_AT_decl_line : 5
    DW_TAG_member
      DW_AT_name : "x"
      DW_AT_decl_file : 1
      DW_AT_decl_line : 6
      DW_AT_type : reference to L2
      DW_AT_data_member_location : 0
    DW_TAG_member
      DW_AT_name : "y"
      DW_AT_decl_file : 1
      DW_AT_decl_line : 7
      DW_AT_type : reference to L2
      DW_AT_data_member_location : 4
L2:
  DW_TAG_base_type
    DW_AT_byte_size : 4
    DW_AT_encoding : DW_ATE_signed
    DW_AT_name : "int"
```

Figure E.10: Type signature computation #1: DWARF representation

1 Running an MD5 hash over this byte stream, and taking the low-order 64 bits,  
2 yields the final signature: 0xd28081e8 dcf5070a.

3 Next, consider a representation of the DWARF information that describes the  
4 type “class A” as shown in Figure E.12 on page 401.

5 In this example, the structure types N : :A and N : :C have each been placed in  
6 separate type units. For N : :A, the actual definition of the type begins at label L1.  
7 The definition involves references to the int base type and to two pointer types.  
8 The information for each of these referenced types is also included in this type  
9 unit, since base types and pointer types are trivial types that are not worth the  
10 overhead of a separate type unit. The last pointer type contains a reference to an  
11 incomplete type N : :B, which is also included here as a declaration, since the  
12 complete type is unknown and its signature is therefore unavailable. There is  
13 also a reference to N : :C, using DW\_FORM\_ref\_sig8 to refer to the type signature  
14 for that type.

## Appendix E. Compression (Informative)

```
// Step 2: 'C' DW_TAG_namespace "N"
0x43 0x39 0x4e 0x00
// Step 3: 'D' DW_TAG_structure_type
0x44 0x13
// Step 4: 'A' DW_AT_name DW_FORM_string "C"
0x41 0x03 0x08 0x43 0x00
// Step 4: 'A' DW_AT_byte_size DW_FORM_sdata 8
0x41 0x0b 0x0d 0x08
// Step 7: First child ("x")
  // Step 3: 'D' DW_TAG_member
  0x44 0x0d
  // Step 4: 'A' DW_AT_name DW_FORM_string "x"
  0x41 0x03 0x08 0x78 0x00
  // Step 4: 'A' DW_AT_data_member_location DW_FORM_sdata 0
  0x41 0x38 0x0d 0x00
  // Step 6: 'T' DW_AT_type (type #2)
  0x54 0x49
    // Step 3: 'D' DW_TAG_base_type
    0x44 0x24
    // Step 4: 'A' DW_AT_name DW_FORM_string "int"
    0x41 0x03 0x08 0x69 0x6e 0x74 0x00
    // Step 4: 'A' DW_AT_byte_size DW_FORM_sdata 4
    0x41 0x0b 0x0d 0x04
    // Step 4: 'A' DW_AT_encoding DW_FORM_sdata DW_ATE_signed
    0x41 0x3e 0x0d 0x05
    // Step 7: End of DW_TAG_base_type "int"
    0x00
  // Step 7: End of DW_TAG_member "x"
  0x00
// Step 7: Second child ("y")
  // Step 3: 'D' DW_TAG_member
  0x44 0x0d
  // Step 4: 'A' DW_AT_name DW_FORM_string "y"
  0x41 0x03 0x08 0x79 0x00
  // Step 4: 'A' DW_AT_data_member_location DW_FORM_sdata 4
  0x41 0x38 0x0d 0x04
  // Step 6: 'R' DW_AT_type (type #2)
  0x52 0x49 0x02
  // Step 7: End of DW_TAG_member "y"
  0x00
// Step 7: End of DW_TAG_structure_type "C"
0x00
```

Figure E.11: Type signature computation #1: flattened byte stream

## Appendix E. Compression (Informative)

part 1 of 2

```
DW_TAG_type_unit
  DW_AT_language_name : DW_LNAME_C_plus_plus (4)
  DW_TAG_namespace
    DW_AT_name : "N"
L1:
  DW_TAG_class_type
    DW_AT_name : "A"
    DW_AT_byte_size : 20
    DW_AT_decl_file : 1
    DW_AT_decl_line : 10
    DW_TAG_member
      DW_AT_name : "v_"
      DW_AT_decl_file : 1
      DW_AT_decl_line : 15
      DW_AT_type : reference to L2
      DW_AT_data_member_location : 0
      DW_AT_accessibility : DW_ACCESS_private
    DW_TAG_member
      DW_AT_name : "next"
      DW_AT_decl_file : 1
      DW_AT_decl_line : 16
      DW_AT_type : reference to L3
      DW_AT_data_member_location : 4
      DW_AT_accessibility : DW_ACCESS_private
    DW_TAG_member
      DW_AT_name : "bp"
      DW_AT_decl_file : 1
      DW_AT_decl_line : 17
      DW_AT_type : reference to L4
      DW_AT_data_member_location : 8
      DW_AT_accessibility : DW_ACCESS_private
    DW_TAG_member
      DW_AT_name : "c"
      DW_AT_decl_file : 1
      DW_AT_decl_line : 18
      DW_AT_type : 0xd28081e8 dcf5070a (signature for struct C)
      DW_AT_data_member_location : 12
      DW_AT_accessibility : DW_ACCESS_private
```

Figure E.12: Type signature computation #2: DWARF representation

## Appendix E. Compression (Informative)

part 2 of 2

```
DW_TAG_subprogram
  DW_AT_external : 1
  DW_AT_name : "A"
  DW_AT_decl_file : 1
  DW_AT_decl_line : 12
  DW_AT_declaration : 1
  DW_TAG_formal_parameter
    DW_AT_type : reference to L3
    DW_AT_artificial : 1
  DW_TAG_formal_parameter
    DW_AT_type : reference to L2
  DW_TAG_subprogram
    DW_AT_external : 1
    DW_AT_name : "v"
    DW_AT_decl_file : 1
    DW_AT_decl_line : 13
    DW_AT_type : reference to L2
    DW_AT_declaration : 1
    DW_TAG_formal_parameter
      DW_AT_type : reference to L3
      DW_AT_artificial : 1
L2:
  DW_TAG_base_type
    DW_AT_byte_size : 4
    DW_AT_encoding : DW_ATE_signed
    DW_AT_name : "int"
L3:
  DW_TAG_pointer_type
    DW_AT_type : reference to L1
L4:
  DW_TAG_pointer_type
    DW_AT_type : reference to L5
  DW_TAG_namespace
    DW_AT_name : "N"
L5:
  DW_TAG_structure_type
    DW_AT_name : "B"
    DW_AT_declaration : 1
```

Figure E.12: Type signature computation #2: DWARF representation (*concluded*)



## Appendix E. Compression (Informative)

part 1 of 3

```
// Step 2: 'C' DW_TAG_namespace "N"
0x43 0x39 0x4e 0x00
// Step 3: 'D' DW_TAG_class_type
0x44 0x02
// Step 4: 'A' DW_AT_name DW_FORM_string "A"
0x41 0x03 0x08 0x41 0x00
// Step 4: 'A' DW_AT_byte_size DW_FORM_sdata 20
0x41 0x0b 0x0d 0x14
// Step 7: First child ("v_")
  // Step 3: 'D' DW_TAG_member
  0x44 0x0d
  // Step 4: 'A' DW_AT_name DW_FORM_string "v_"
  0x41 0x03 0x08 0x76 0x5f 0x00
  // Step 4: 'A' DW_AT_accessibility DW_FORM_sdata DW_ACCESS_private
  0x41 0x32 0x0d 0x03
  // Step 4: 'A' DW_AT_data_member_location DW_FORM_sdata 0
  0x41 0x38 0x0d 0x00
  // Step 6: 'T' DW_AT_type (type #2)
  0x54 0x49
    // Step 3: 'D' DW_TAG_base_type
    0x44 0x24
    // Step 4: 'A' DW_AT_name DW_FORM_string "int"
    0x41 0x03 0x08 0x69 0x6e 0x74 0x00
    // Step 4: 'A' DW_AT_byte_size DW_FORM_sdata 4
    0x41 0x0b 0x0d 0x04
    // Step 4: 'A' DW_AT_encoding DW_FORM_sdata DW_ATE_signed
    0x41 0x3e 0x0d 0x05
    // Step 7: End of DW_TAG_base_type "int"
    0x00
  // Step 7: End of DW_TAG_member "v_"
  0x00
// Step 7: Second child ("next")
  // Step 3: 'D' DW_TAG_member
  0x44 0x0d
  // Step 4: 'A' DW_AT_name DW_FORM_string "next"
  0x41 0x03 0x08 0x6e 0x65 0x78 0x74 0x00
  // Step 4: 'A' DW_AT_accessibility DW_FORM_sdata DW_ACCESS_private
  0x41 0x32 0x0d 0x03
  // Step 4: 'A' DW_AT_data_member_location DW_FORM_sdata 4
  0x41 0x38 0x0d 0x04
```

Figure E.13: Type signature example #2: flattened byte stream

## Appendix E. Compression (Informative)

part 2 of 3

```
// Step 6: 'T' DW_AT_type (type #3)
0x54 0x49
  // Step 3: 'D' DW_TAG_pointer_type
  0x44 0x0f
  // Step 5: 'N' DW_AT_type
  0x4e 0x49
  // Step 5: 'C' DW_TAG_namespace "N" 'E'
  0x43 0x39 0x4e 0x00 0x45
  // Step 5: "A"
  0x41 0x00
  // Step 7: End of DW_TAG_pointer_type
  0x00
// Step 7: End of DW_TAG_member "next"
0x00
// Step 7: Third child ("bp")
  // Step 3: 'D' DW_TAG_member
  0x44 0x0d
  // Step 4: 'A' DW_AT_name DW_FORM_string "bp"
  0x41 0x03 0x08 0x62 0x70 0x00
  // Step 4: 'A' DW_AT_accessibility DW_FORM_sdata DW_ACCESS_private
  0x41 0x32 0x0d 0x03
  // Step 4: 'A' DW_AT_data_member_location DW_FORM_sdata 8
  0x41 0x38 0x0d 0x08
  // Step 6: 'T' DW_AT_type (type #4)
  0x54 0x49
    // Step 3: 'D' DW_TAG_pointer_type
    0x44 0x0f
    // Step 5: 'N' DW_AT_type
    0x4e 0x49
    // Step 5: 'C' DW_TAG_namespace "N" 'E'
    0x43 0x39 0x4e 0x00 0x45
    // Step 5: "B"
    0x42 0x00
    // Step 7: End of DW_TAG_pointer_type
    0x00
  // Step 7: End of DW_TAG_member "next"
  0x00
// Step 7: Fourth child ("c")
  // Step 3: 'D' DW_TAG_member
  0x44 0x0d
  // Step 4: 'A' DW_AT_name DW_FORM_string "c"
  0x41 0x03 0x08 0x63 0x00
  // Step 4: 'A' DW_AT_accessibility DW_FORM_sdata DW_ACCESS_private
  0x41 0x32 0x0d 0x03
```

Figure E.13: Type signature example #2: flattened byte stream (*continued*)

## Appendix E. Compression (Informative)

part 3 of 3

```
// Step 4: 'A' DW_AT_data_member_location DW_FORM_sdata 12
0x41 0x38 0x0d 0x0c
// Step 6: 'T' DW_AT_type (type #5)
0x54 0x49
  // Step 2: 'C' DW_TAG_namespace "N"
  0x43 0x39 0x4e 0x00
  // Step 3: 'D' DW_TAG_structure_type
  0x44 0x13
  // Step 4: 'A' DW_AT_name DW_FORM_string "C"
  0x41 0x03 0x08 0x43 0x00
  // Step 4: 'A' DW_AT_byte_size DW_FORM_sdata 8
  0x41 0x0b 0x0d 0x08
  // Step 7: First child ("x")
    // Step 3: 'D' DW_TAG_member
    0x44 0x0d
    // Step 4: 'A' DW_AT_name DW_FORM_string "x"
    0x41 0x03 0x08 0x78 0x00
    // Step 4: 'A' DW_AT_data_member_location DW_FORM_sdata 0
    0x41 0x38 0x0d 0x00
    // Step 6: 'R' DW_AT_type (type #2)
    0x52 0x49 0x02
    // Step 7: End of DW_TAG_member "x"
    0x00
  // Step 7: Second child ("y")
    // Step 3: 'D' DW_TAG_member
    0x44 0x0d
    // Step 4: 'A' DW_AT_name DW_FORM_string "y"
    0x41 0x03 0x08 0x79 0x00
    // Step 4: 'A' DW_AT_data_member_location DW_FORM_sdata 4
    0x41 0x38 0x0d 0x04
    // Step 6: 'R' DW_AT_type (type #2)
    0x52 0x49 0x02
    // Step 7: End of DW_TAG_member "y"
    0x00
  // Step 7: End of DW_TAG_structure_type "C"
  0x00
// Step 7: End of DW_TAG_member "c"
0x00
// Step 7: Fifth child ("A")
  // Step 3: 'S' DW_TAG_subprogram "A"
  0x53 0x2e 0x41 0x00
// Step 7: Sixth child ("v")
  // Step 3: 'S' DW_TAG_subprogram "v"
  0x53 0x2e 0x76 0x00
// Step 7: End of DW_TAG_structure_type "A"
0x00
```

Figure E.13: Type signature example #2: flattened byte stream (*concluded*)

## Appendix E. Compression (Informative)

1 In computing a signature for the type `N :: A`, flatten the type description into a  
2 byte stream according to the procedure outlined in Section 7.32 on page 257. The  
3 result is shown in Figure E.13 on page 403.

4 Running an MD5 hash over this byte stream, and taking the low-order 64 bits,  
5 yields the final signature: 0xd6d160f5 5589f6e9.

6 A source file that includes this header file may declare a variable of type `N :: A`,  
7 and its DWARF information may look like that shown in Figure E.14.

```
DW_TAG_compile_unit
...
DW_TAG_subprogram
...
  DW_TAG_variable
    DW_AT_name : "a"
    DW_AT_type : (signature) 0xd6d160f5 5589f6e9
    DW_AT_location : ...
  ...
```

Figure E.14: Type signature example usage

### 8 E.2.2 Type Signature Computation Grammar

9 Figure E.15 on the next page presents a semi-formal grammar that may aid in  
10 understanding how the bytes of the flattened type description are formed during  
11 the type signature computation algorithm of Section 7.32 on page 257.

## Appendix E. Compression (Informative)

```
signature
  : opt-context debug-entry attributes children
opt-context          // Step 2
  : 'C' tag-code string opt-context
  : empty
debug-entry          // Step 3
  : 'D' tag-code
attributes           // Steps 4, 5, 6
  : attribute attributes
  : empty
attribute
  : 'A' at-code form-encoded-value // Normal attributes
  : 'N' at-code opt-context 'E' string // Reference to type by name
  : 'R' at-code back-ref           // Back-reference to visited type
  : 'T' at-code signature          // Recursive type
children             // Step 7
  : child children
  : '\0'
child
  : 'S' tag-code string
  : signature
tag-code
  : <ULEB128>
at-code
  : <ULEB128>
form-encoded-value
  : DW_FORM_sdata value
  : DW_FORM_flag value
  : DW_FORM_string string
  : DW_FORM_block block
DW_FORM_string
  : '\x08'
DW_FORM_block
  : '\x09'
DW_FORM_flag
  : '\x0c'
DW_FORM_sdata
  : '\x0d'
value
  : <SLEB128>
block
  : <ULEB128> <fixed-length-block> // The ULEB128 gives the length of the block
back-ref
  : <ULEB128>
string
  : <null-terminated-string>
empty
  :
  :
```

Figure E.15: Type signature computation grammar

### E.2.3 Declarations Completing Non-Defining Declarations

Consider a compilation unit that contains a definition of the member function `N::A::v()` from Figure E.9 on page 398. A possible representation of the debug information for this function in the compilation unit is shown in Figure E.16.

```

DW_TAG_namespace
  DW_AT_name : "N"
L1:
  DW_TAG_class_type
    DW_AT_name : "A"
    DW_AT_declaration : true
    DW_AT_signature : 0xd6d160f5 5589f6e9
L2:
  DW_TAG_subprogram
    DW_AT_external : 1
    DW_AT_name : "v"
    DW_AT_decl_file : 1
    DW_AT_decl_line : 13
    DW_AT_type : reference to L3
    DW_AT_declaration : 1
    DW_TAG_formal_parameter
      DW_AT_type : reference to L4
      DW_AT_artificial : 1
...
L3:
  DW_TAG_base_type
    DW_AT_byte_size : 4
    DW_AT_encoding : DW_ATE_signed
    DW_AT_name : "int"
...
L4:
  DW_TAG_pointer_type
    DW_AT_type : reference to L1
...
  DW_TAG_subprogram
    DW_AT_specification : reference to L2
    DW_AT_decl_file : 2
    DW_AT_decl_line : 25
    DW_AT_low_pc : ...
    DW_AT_high_pc : ...
  DW_TAG_lexical_block
    ...
...

```

Figure E.16: Completing declaration of a member function: DWARF encoding

## 1 **E.3 Summary of Compression Techniques**

### 2 **E.3.1 #include compression**

3 C++ has a much greater problem than C with the number and size of the headers  
4 included and the amount of data in each, but even with C there is substantial  
5 header file information duplication.

6 A reasonable approach is to put each header file in its own section group, using  
7 the naming rules mentioned above. The section groups are marked to ensure  
8 duplicate removal.

9 All data instances and code instances (even if they came from the header files  
10 above) are put into non-section group sections such as the base object file  
11 `.debug_info` section.

### 12 **E.3.2 Eliminating function duplication**

13 Function templates (C++) result in code for the same template instantiation being  
14 compiled into multiple archives or relocatable object files. The linker wants to  
15 keep only one of a given entity. The DWARF description, and everything else for  
16 this function, should be reduced to just a single copy.

17 For each such code group (function template in this example) the compiler  
18 assigns a name for the group which will match all other instantiations of this  
19 function but match nothing else. The section groups are marked to ensure  
20 duplicate removal, so that the second and subsequent definitions seen by the  
21 static linker are simply discarded.

22 References to other `.debug_info` sections follow the approach suggested above,  
23 but the naming rule is slightly different in that the `<file-designator>` should be  
24 interpreted as a `<function-designator>`.

### 25 **E.3.3 Single-function-per-DWARF-compilation-unit**

26 Section groups can help make it easy for a linker to completely remove unused  
27 functions.

28 Such section groups are not marked for duplicate removal, since the functions  
29 are not duplicates of anything.

30 Each function is given a compilation unit and a section group. Each such  
31 compilation unit is complete, with its own text, data, and DWARF sections.

## Appendix E. Compression (Informative)

1 There will also be a compilation unit that has the file-level declarations and  
2 definitions. Other per-function compilation unit DWARF information  
3 (`.debug_info`) points to this common file-level compilation unit using  
4 `DW_TAG_imported_unit`.

5 Section groups can use `DW_FORM_ref_addr` and internal labels (section-relative  
6 relocations) to refer to the main object file sections, as the section groups here are  
7 either deleted as unused or kept. There is no possibility (aside from error) of a  
8 group from some other compilation being used in place of one of these groups.

### 9 **E.3.4 Inlining and out-of-line-instances**

10 Abstract instances and concrete-out-of-line instances may be put in distinct  
11 compilation units using section groups. This makes possible some useful  
12 duplicate DWARF elimination.

13 *No special provision for eliminating class duplication resulting from template*  
14 *instantiation is made here, though nothing prevents eliminating such duplicates using*  
15 *section groups.*

### 16 **E.3.5 Separate Type Units**

17 Each complete declaration of a globally-visible type can be placed in its own  
18 separate type section, with a group key derived from the type signature. The  
19 linker can then remove all duplicate type declarations based on the key.



Appendix E. Compression (Informative)

*(empty page)*

# Appendix F

## Split DWARF Object Files (Informative)

With the traditional DWARF format, debug information is designed with the expectation that it will be processed by the linker to produce an output binary with complete debug information, and with fully-resolved references to locations within the application. For very large applications, however, this approach can result in excessively large link times and excessively large output files.

Several vendors have independently developed proprietary approaches that allow the debug information to remain in the relocatable object files, so that the linker does not have to process the debug information or copy it to the output file. These approaches have all required that additional information be made available to the debug information consumer, and that the consumer perform some minimal amount of relocation in order to interpret the debug info correctly. The additional information required, in the form of load maps or symbol tables, and the details of the relocation are not covered by the DWARF specification, and vary with each vendor's implementation.

Section [7.3.2 on page 195](#) describes a platform-independent mechanism that allows a producer to split the debugging information into relocatable and non-relocatable partitions. This Appendix describes the use of split DWARF object files and provides some illustrative examples.

### F.1 Overview

DWARF Version 5 introduces an optional set of debugging sections that allow the compiler to partition the debugging information into a set of (small) sections that require link-time relocation and a set of (large) sections that do not. The sections

## Appendix F. Split DWARF Object Files (Informative)

1 that require relocation are written to the relocatable object file as usual, and are  
2 linked into the final executable. The sections that do not require relocation,  
3 however, can be written to the relocatable object (.o) file but ignored by the  
4 linker, or they can be written to a separate DWARF object (.dwo) file that need  
5 not be accessed by the linker.

6 The optional set of debugging sections includes the following:

- 7 • `.debug_abbrev.dwo` - Contains the abbreviations table(s) used by the  
8 `.debug_info.dwo` section.
- 9 • `.debug_info.dwo` - Contains the [DW\\_TAG\\_compile\\_unit](#) and  
10 [DW\\_TAG\\_type\\_unit](#) DIEs and their descendants. This is the bulk of the  
11 debugging information for the compilation unit that is normally found in  
12 the `.debug_info` section.
- 13 • `.debug_loclists.dwo` - Contains the location lists referenced by the  
14 debugging information entries in the `.debug_info.dwo` section. This  
15 contains the location lists normally found in the `.debug_loclists` section.
- 16 • `.debug_str.dwo` - Contains the string table for all indirect strings  
17 referenced by the debugging information in the `.debug_info.dwo` sections.
- 18 • `.debug_str_offsets.dwo` - Contains the string offsets table for the strings  
19 in the `.debug_str.dwo` section.
- 20 • `.debug_macro.dwo` - Contains macro definition information, normally  
21 found in the `.debug_macro` section.
- 22 • `.debug_line.dwo` - Contains specialized line number tables for the type  
23 units in the `.debug_info.dwo` section. These tables contain only the  
24 directory and filename lists needed to interpret [DW\\_AT\\_decl\\_file](#) attributes  
25 in the debugging information entries. Actual line number tables remain in  
26 the `.debug_line` section, and remain in the relocatable object (.o) files.

27 In a `.dwo` file, there is no benefit to having a separate string section for directories  
28 and file names because the primary string table will never be stripped.

29 Accordingly, no `.debug_line_str.dwo` section is defined. Content descriptions  
30 corresponding to [DW\\_FORM\\_line\\_strp](#) in an executable file (for example, in the  
31 skeleton compilation unit) instead use one of the forms [DW\\_FORM\\_strx](#),  
32 [DW\\_FORM\\_strx1](#), [DW\\_FORM\\_strx2](#), [DW\\_FORM\\_strx3](#) or [DW\\_FORM\\_strx4](#).

33 This allows directory and file name strings to be merged with general strings and  
34 across compilations in package files (where they are not subject to potential  
35 stripping). This merge is facilitated by the requirement that all references to the  
36 `.debug_str.dwo` string table are made indirectly through the

## Appendix F. Split DWARF Object Files (Informative)

1 .debug\_str\_offsets.dwo section so that only that section needs to be modified  
2 during string merging (see Section 7.3.2.2 on page 196).

3 In order for the consumer to locate and process the debug information, the  
4 compiler must produce a small amount of debug information that passes through  
5 the linker into the output binary. A skeleton .debug\_info section for each  
6 compilation unit contains a reference to the corresponding .o or .dwo file, and  
7 the .debug\_line section (which is typically small compared to the .debug\_info  
8 sections) is linked into the output binary, as is the .debug\_addr section.

9 The debug sections that continue to be linked into the output binary include the  
10 following:

- 11 • .debug\_abbrev - Contains the abbreviation codes used by the skeleton  
12 .debug\_info section.
- 13 • .debug\_addr - Contains references to loadable sections, indexed by  
14 attributes of one of the forms DW\_FORM\_addrx, DW\_FORM\_addrx1,  
15 DW\_FORM\_addrx2, DW\_FORM\_addrx3, DW\_FORM\_addrx4, or location  
16 expression DW\_OP\_addrx opcodes.
- 17 • .debug\_aranges - Contains the accelerated range lookup table for the  
18 compilation unit.
- 19 • .debug\_frame - Contains the frame tables.
- 20 • .debug\_info - Contains a skeleton compilation unit DIE, which  
21 has no children.
- 22 • .debug\_line - Contains the line number tables. (These could be moved to  
23 the .dwo file, but in order to do so, each DW\_LNE\_set\_address opcode  
24 would need to be replaced by a new opcode that referenced an entry in the  
25 .debug\_addr section. Furthermore, leaving this section in the .o file allows  
26 many debug info consumers to remain unaware of .dwo files.)
- 27 • .debug\_line\_str - Contains strings for file names used in combination  
28 with the .debug\_line section.
- 29 • .debug\_names - Contains the names for use in building an index section.  
30 The section header refers to a compilation unit offset, which is the offset of  
31 the skeleton compilation unit in the .debug\_info section.
- 32 • .debug\_str - Contains any strings referenced by the skeleton .debug\_info  
33 sections (via DW\_FORM\_strp, DW\_FORM\_strp8, DW\_FORM\_strx,  
34 DW\_FORM\_strx1, DW\_FORM\_strx2, DW\_FORM\_strx3 or  
35 DW\_FORM\_strx4).

## Appendix F. Split DWARF Object Files (Informative)

- 1 • `.debug_str_offsets` - Contains the string offsets table for the strings in the  
2 `.debug_str` section (if one of the forms `DW_FORM_strx`, `DW_FORM_strx1`,  
3 `DW_FORM_strx2`, `DW_FORM_strx3` or `DW_FORM_strx4` is used).

4 The skeleton compilation unit DIE may have the following attributes:

`DW_AT_addr_base` `DW_AT_high_pc` `DW_AT_stmt_list`  
`DW_AT_comp_dir` `DW_AT_low_pc` `DW_AT_str_offsets`  
`DW_AT_dwo_name` `DW_AT_ranges`

5 All other attributes of the compilation unit DIE are moved to the full DIE in the  
6 `.debug_info.dwo` section.

7 The `dwo_id` field is present in headers of the skeleton DIE and the header of the  
8 full DIE, so that a consumer can verify a match.

9 Relocations are neither necessary nor useful in `.dwo` files, because the `.dwo` files  
10 contain only debugging information that does not need to be processed by a  
11 linker. Relocations are rendered unnecessary by these strategies:

- 12 1. Some values needing relocation are kept in the `.o` file (for example, references  
13 to the line number program from the skeleton compilation unit).
- 14 2. Some values do not need a relocation because they refer from one `.dwo`  
15 section to another `.dwo` section in the same compilation unit.
- 16 3. Some values that need a relocation to refer to a relocatable program address  
17 use one of the `DW_FORM_addrx`, `DW_FORM_addrx1`, `DW_FORM_addrx2`,  
18 `DW_FORM_addrx3` or `DW_FORM_addrx4` forms, referencing a relocatable  
19 value in the `.debug_addr` section (which remains in the `.o` file).

20 Table F.1 on the next page summarizes which attributes are defined for use in the  
21 various kinds of compilation units (see Section 3.1 on page 61). It compares and  
22 contrasts both conventional and split object-related kinds.

23 The split dwarf object file design depends on having an index of debugging  
24 information available to the consumer. For name lookups, the consumer can use  
25 the `.debug_names` index section (see Section 6.1 on page 141) to locate a skeleton  
26 compilation unit. The `DW_AT_comp_dir` and `DW_AT_dwo_name` attributes in  
27 the skeleton compilation unit can then be used to locate the corresponding  
28 DWARF object file for the compilation unit. Similarly, for an address lookup, the  
29 consumer can use the `.debug_aranges` table, which will also lead to a skeleton  
30 compilation unit. For a file and line number lookup, the skeleton compilation  
31 units can be used to locate the line number tables.

## Appendix F. Split DWARF Object Files (Informative)

Table F.1: Unit attributes by unit kind

Attribute	Unit Kind				
	Conventional Full & Partial	Type	Skeleton	Split Full	Split Type
DW_AT_addr_base	✓		✓		
DW_AT_base_types	✓				
DW_AT_comp_dir	✓		✓		
DW_AT_dwo_name			✓		
DW_AT_entry_pc	✓			✓	
DW_AT_high_pc	✓		✓		
DW_AT_identifier_case	✓			✓	
DW_AT_language_name	✓	✓		✓	✓
DW_AT_language_version	✓	✓		✓	✓
DW_AT_loclists_base	✓				
DW_AT_low_pc	✓		✓		
DW_AT_macros	✓			✓	
DW_AT_main_subprogram	✓			✓	
DW_AT_name	✓			✓	
DW_AT_producer	✓			✓	
DW_AT_ranges	✓			✓	
DW_AT_rnglists_base	✓		✓		
DW_AT_stmt_list	✓	✓	✓		✓
DW_AT_str_offsets	✓	✓	✓		
DW_AT_use_UTF8	✓	✓	✓	✓	✓

## F.2 Split DWARF Object File Example

Consider the example source code in Figure F.1, Figure F.2 on the next page and Figure F.3 on page 419. When compiled with split DWARF, we will have two DWARF object files, `demo1.o` and `demo2.o`, and two split DWARF object files, `demo1.dwo` and `demo2.dwo`.

In this section, we will use this example to show how the connections between the relocatable object file and the split DWARF object file are maintained through the linking process. In the next section, we will use this same example to show how two or more split DWARF object files are combined into a DWARF package file.

*File demo1.cc*

```
#include "demo.h"

bool Box::contains(const Point& p) const
{
    return (p.x() >= ll_.x() && p.x() <= ur_.x() &&
            p.y() >= ll_.y() && p.y() <= ur_.y());
}
```

Figure F.1: Split object example: source fragment #1

## Appendix F. Split DWARF Object Files (Informative)

*File demo2.cc*

```
#include "demo.h"

bool Line::clip(const Box& b)
{
    float slope = (end_.y() - start_.y()) / (end_.x() - start_.x());
    while (1) {
        // Trivial acceptance.
        if (b.contains(start_) && b.contains(end_)) return true;

        // Trivial rejection.
        if (start_.x() < b.l() && end_.x() < b.l()) return false;
        if (start_.x() > b.r() && end_.x() > b.r()) return false;
        if (start_.y() < b.b() && end_.y() < b.b()) return false;
        if (start_.y() > b.t() && end_.y() > b.t()) return false;

        if (b.contains(start_)) {
            // Swap points so that start_ is outside the clipping
            // rectangle.
            Point temp = start_;
            start_ = end_;
            end_ = temp;
        }

        if (start_.x() < b.l())
            start_ = Point(b.l(),
                          start_.y() + (b.l() - start_.x()) * slope);
        else if (start_.x() > b.r())
            start_ = Point(b.r(),
                          start_.y() + (b.r() - start_.x()) * slope);
        else if (start_.y() < b.b())
            start_ = Point(start_.x() + (b.b() - start_.y()) / slope,
                          b.b());
        else if (start_.y() > b.t())
            start_ = Point(start_.x() + (b.t() - start_.y()) / slope,
                          b.t());
    }
}
```

Figure F.2: Split object example: source fragment #2



## Appendix F. Split DWARF Object Files (Informative)

*File demo.h*

```
class A {
public:
    Point(float x, float y) : x_(x), y_(y){}
    float x() const { return x_; }
    float y() const { return y_; }
private:
    float x_;
    float y_;
};

class Line {
public:
    Line(Point start, Point end) : start_(start), end_(end){}
    bool clip(const Box& b);
    Point start() const { return start_; }
    Point end() const { return end_; }
private:
    Point start_;
    Point end_;
};

class Box {
public:
    Box(float l, float r, float b, float t) : ll_(l, b), ur_(r, t){}
    Box(Point ll, Point ur) : ll_(ll), ur_(ur){}
    bool contains(const Point& p) const;
    float l() const { return ll_.x(); }
    float r() const { return ur_.x(); }
    float b() const { return ll_.y(); }
    float t() const { return ur_.y(); }
private:
    Point ll_;
    Point ur_;
};
```

Figure F.3: Split object example: source fragment #3

## 1 F.2.1 Contents of the Object Files

2 The object files each contain the following sections of debug information:

```
3     .debug_abbrev
4     .debug_info
5     .debug_line
6     .debug_str
7     .debug_addr
8     .debug_names
9     .debug_aranges
```

10 The `.debug_abbrev` section contains just a single entry describing the skeleton  
11 compilation unit DIE.

12 The DWARF description in the `.debug_info` section contains just a single DIE,  
13 the skeleton compilation unit, which may look like Figure F.4 following.

```
DW_TAG_skeleton_unit
  DW_AT_comp_dir: (reference to directory name in .debug_str)
  DW_AT_dwo_name: (reference to "demo1.dwo" in .debug_str)
  DW_AT_addr_base: (reference to .debug_addr section)
  DW_AT_stmt_list: (reference to .debug_line section)
```

Figure F.4: Split object example: skeleton DWARF description

14 The `DW_AT_comp_dir` and `DW_AT_dwo_name` attributes provide the location  
15 of the corresponding split DWARF object file that contains the full debug  
16 information; that file is normally expected to be in the same directory as the  
17 object file itself.

18 The `dwo_id` field in the header of the skeleton unit provides an ID or key for the  
19 debug information contained in the DWARF object file. This ID serves two  
20 purposes: it can be used to verify that the debug information in the split DWARF  
21 object file matches the information in the object file, and it can be used to find the  
22 debug information in a DWARF package file.

23 The `DW_AT_addr_base` attribute contains the relocatable offset of this object  
24 file's contribution to the `.debug_addr` section.

25 The `DW_AT_stmt_list` attribute contains the relocatable offset of this file's  
26 contribution to the `.debug_line` table.

## Appendix F. Split DWARF Object Files (Informative)

1 The `.debug_line` section contains the full line number table for the compiled  
2 code in the object file. As shown in Figure F.1 on page 417, the line number  
3 program header lists the two file names, `demo.h` and `demo1.cc`, and contains line  
4 number programs for `Box::contains`, `Point::x`, and `Point::y`.

5 The `.debug_str` section contains the strings referenced indirectly by the  
6 compilation unit DIE and by the line number program.

7 The `.debug_addr` section contains relocatable addresses of locations in the  
8 loadable text and data that are referenced by debugging information entries in  
9 the split DWARF object. In the example in F.3 on page 419, `demo1.o` may have  
10 three entries:

Slot	Location referenced
0	low PC value for <code>Box::contains</code>
1	low PC value for <code>Point::x</code>
2	low PC value for <code>Point::y</code>

11 The `.debug_names` section contains the names defined by the debugging  
12 information in the split DWARF object file (see Section 6.1.1.1 on page 143), and  
13 references the skeleton compilation unit. When linked together into a final  
14 executable, they can be used by a DWARF consumer to lookup a name to find  
15 one or more skeleton compilation units that provide information about that  
16 name. From the skeleton compilation unit, the consumer can find the split  
17 DWARF object file that it can then read to get the full DWARF information.

18 The `.debug_aranges` section contains the PC ranges defined in this compilation  
19 unit, and allow a DWARF consumer to map a PC value to a skeleton compilation  
20 unit, and then to a split DWARF object file.

### 21 F.2.2 Contents of the Linked Executable File

22 When `demo1.o` and `demo2.o` are linked together (along with a main program and  
23 other necessary library routines that we will ignore here for simplicity), the  
24 resulting executable file will contain at least the two skeleton compilation units  
25 in the `.debug_info` section, as shown in Figure F.5 following.

26 Each skeleton compilation unit has a `DW_AT_stmt_list` attribute, which provides  
27 the relocated offset to that compilation unit's contribution in the executable's  
28 `.debug_line` section. In this example, the line number information for `demo1.dwo`  
29 begins at offset 120, and for `demo2.dwo`, it begins at offset 200.

## Appendix F. Split DWARF Object Files (Informative)

```
DW_TAG_skeleton_unit
  DW_AT_comp_dir: (reference to directory name in .debug_str)
  DW_AT_dwo_name: (reference to "demo1.dwo" in .debug_str)
  DW_AT_addr_base: 48 (offset in .debug_addr)
  DW_AT_stmt_list: 120 (offset in .debug_line)
DW_TAG_skeleton_unit
  DW_AT_comp_dir: (reference to directory name in .debug_str)
  DW_AT_dwo_name: (reference to "demo2.dwo" in .debug_str)
  DW_AT_addr_base: 80 (offset in .debug_addr)
  DW_AT_stmt_list: 200 (offset in .debug_line)
```

Figure F.5: Split object example: executable file DWARF excerpts

1 Each skeleton compilation unit also has a `DW_AT_addr_base` attribute, which  
2 provides the relocated offset to that compilation unit's contribution in the  
3 executable's `.debug_addr` section. Unlike the `DW_AT_stmt_list` attribute, the  
4 offset refers to the first address table slot, not to the section header. In this  
5 example, we see that the first address (slot 0) from `demo1.o` begins at offset 48.  
6 Because the `.debug_addr` section contains an 8-byte header, the object file's  
7 contribution to the section actually begins at offset 40 (for a 64-bit DWARF object,  
8 the header would be 16 bytes long, and the value for the `DW_AT_addr_base`  
9 attribute would then be 56). All attributes in `demo1.dwo` that use  
10 `DW_FORM_addrx`, `DW_FORM_addrx1`, `DW_FORM_addrx2`,  
11 `DW_FORM_addrx3` or `DW_FORM_addrx4` would then refer to address table  
12 slots relative to that offset. Likewise, the `.debug_addr` contribution from  
13 `demo2.dwo` begins at offset 72, and its first address slot is at offset 80. Because  
14 these contributions have been processed by the linker, they contain relocated  
15 values for the addresses in the program that are referred to by the debug  
16 information.

17 The linked executable will also contain `.debug_abbrev`, `.debug_str`,  
18 `.debug_names` and `.debug_aranges` sections, each the result of combining and  
19 relocating the contributions from the relocatable object files.

## F.2.3 Contents of the Split DWARF Object Files

The split DWARF object files each contain the following sections:

```

3     .debug_abbrev.dwo
4     .debug_info.dwo (for the compilation unit)
5     .debug_info.dwo (one COMDAT section for each type unit)
6     .debug_loclists.dwo
7     .debug_line.dwo
8     .debug_macro.dwo
9     .debug_rnglists.dwo
10    .debug_str_offsets.dwo
11    .debug_str.dwo

```

The `.debug_abbrev.dwo` section contains the abbreviation declarations for the debugging information entries in the `.debug_info.dwo` section.

The `.debug_info.dwo` section containing the compilation unit contains the full debugging information for the compile unit, and looks much like a normal `.debug_info` section in a non-split object file, with the following exceptions:

- The `DW_TAG_compile_unit` DIE does not need to repeat the `DW_AT_ranges`, `DW_AT_low_pc`, `DW_AT_high_pc`, and `DW_AT_stmt_list` attributes that are provided in the skeleton compilation unit.
- References to strings in the string table use the form code `DW_FORM_strx`, `DW_FORM_strx1`, `DW_FORM_strx2`, `DW_FORM_strx3` or `DW_FORM_strx4`, referring to slots in the `.debug_str_offsets.dwo` section.
- References to relocatable addresses in the object file use one of the form codes `DW_FORM_addrx`, `DW_FORM_addrx1`, `DW_FORM_addrx2`, `DW_FORM_addrx3` or `DW_FORM_addrx4`, referring to slots in the `.debug_addr` table, relative to the base offset given by `DW_AT_addr_base` in the skeleton compilation unit.

Figure F.6 following presents excerpts from the `.debug_info.dwo` section for `demo1.dwo`.

In the defining declaration for `Box::contains` at 5\$, the `DW_AT_low_pc` attribute is represented using `DW_FORM_addrx`, which refers to slot 0 in the `.debug_addr` table from `demo1.o`. That slot contains the relocated address of the beginning of the function.

## Appendix F. Split DWARF Object Files (Informative)

part 1 of 2

```
DW_TAG_compile_unit
  DW_AT_producer [DW_FORM_strx]: (slot 15) (producer string)
  DW_AT_language_name: DW_LNAME_C_plus_plus
  DW_AT_name [DW_FORM_strx]: (slot 7) "demo1.cc"
  DW_AT_comp_dir [DW_FORM_strx]: (slot 4) (directory name)
1$: DW_TAG_class_type
    DW_AT_name [DW_FORM_strx]: (slot 12) "Point"
    DW_AT_signature [DW_FORM_ref_sig8]: 0x2f33248f03ff18ab
    DW_AT_declaration: true
2$: DW_TAG_subprogram
    DW_AT_external: true
    DW_AT_name [DW_FORM_strx]: (slot 12) "Point"
    DW_AT_decl_file: 1
    DW_AT_decl_line: 5
    DW_AT_linkage_name [DW_FORM_strx]: (slot 16) "_ZN5PointC4Eff"
    DW_AT_accessibility: DW_ACCESS_public
    DW_AT_declaration: true
    ...
3$: DW_TAG_class_type
    DW_AT_name [DW_FORM_string]: "Box"
    DW_AT_signature [DW_FORM_ref_sig8]: 0xe97a3917c5a6529b
    DW_AT_declaration: true
    ...
4$: DW_TAG_subprogram
    DW_AT_external: true
    DW_AT_name [DW_FORM_strx]: (slot 0) "contains"
    DW_AT_decl_file: 1
    DW_AT_decl_line: 28
    DW_AT_linkage_name [DW_FORM_strx]: (slot 8)
                                         "_ZNK3Box8containsERK5Point"
    DW_AT_type: (reference to 7$)
    DW_AT_accessibility: DW_ACCESS_public
    DW_AT_declaration: true
    ...
```

Figure F.6: Split object example: demo1.dwo excerpts

1 Each type unit is contained in its own COMDAT `.debug_info.dwo` section, and  
2 looks like a normal type unit in a non-split object, except that the  
3 `DW_TAG_type_unit` DIE contains a `DW_AT_stmt_list` attribute that refers to a  
4 specialized `.debug_line.dwo` section. This section contains a normal line  
5 number program header with a list of include directories and filenames, but no  
6 line number program. This section is used only as a reference for filenames  
7 needed for `DW_AT_decl_file` attributes within the type unit.

## Appendix F. Split DWARF Object Files (Informative)

part 2 of 2

```
5$: DW_TAG_subprogram
    DW_AT_specification: (reference to 4$)
    DW_AT_decl_file: 2
    DW_AT_decl_line: 3
    DW_AT_low_pc [DW_FORM_addrx]: (slot 0)
    DW_AT_high_pc [DW_FORM_data8]: 0xbb
    DW_AT_frame_base: DW_OP_call_frame_cfa
    DW_AT_object_pointer: (reference to 6$)
6$: DW_TAG_formal_parameter
    DW_AT_name [DW_FORM_strx]: (slot 13): "this"
    DW_AT_type: (reference to 8$)
    DW_AT_artificial: true
    DW_AT_location: DW_OP_fbreg(-24)
    DW_TAG_formal_parameter
    DW_AT_name [DW_FORM_string]: "p"
    DW_AT_decl_file: 2
    DW_AT_decl_line: 3
    DW_AT_type: (reference to 11$)
    DW_AT_location: DW_OP_fbreg(-32)
...
7$: DW_TAG_base_type
    DW_AT_byte_size: 1
    DW_AT_encoding: DW_ATE_boolean
    DW_AT_name [DW_FORM_strx]: (slot 5) "bool"
...
8$: DW_TAG_const_type
    DW_AT_type: (reference to 9$)
9$: DW_TAG_pointer_type
    DW_AT_byte_size: 8
    DW_AT_type: (reference to 10$)
10$: DW_TAG_const_type
    DW_AT_type: (reference to 3$)
...
11$: DW_TAG_const_type
    DW_AT_type: (reference to 12$)
12$: DW_TAG_reference_type
    DW_AT_byte_size: 8
    DW_AT_type: (reference to 13$)
13$: DW_TAG_const_type
    DW_AT_type: (reference to 1$)
...
```

Figure F.6: Split object example: demo1.dwo DWARF excerpts (*concluded*)

## Appendix F. Split DWARF Object Files (Informative)

1 The `.debug_str_offsets.dwo` section contains an entry for each unique string in  
2 the string table. Each entry in the table is the offset of the string, which is  
3 contained in the `.debug_str.dwo` section.

4 In a split DWARF object file, all references to strings go through this table (there  
5 are no other offsets to `.debug_str.dwo` in a split DWARF object file). That is,  
6 there is no use of [DW\\_FORM\\_strp](#) in a split DWARF object file.

7 The offsets in these slots have no associated relocations, because they are not part  
8 of a relocatable object file. When combined into a DWARF package file, however,  
9 each slot must be adjusted to refer to the appropriate offset within the merged  
10 string table (`.debug_str.dwo`). The tool that builds the DWARF package file must  
11 understand the structure of the `.debug_str_offsets.dwo` section in order to  
12 apply the necessary adjustments. [Section F.3 on page 430](#) presents an example of  
13 a DWARF package file.

14 The `.debug_rnglists.dwo` section contains range lists referenced by any  
15 [DW\\_AT\\_ranges](#) attributes in the split DWARF object. In our example, `demo1.o`  
16 would have just a single range list for the compilation unit, with range list entries  
17 for the function `Box::contains` and for out-of-line copies of the inline functions  
18 `Point::x` and `Point::y`.

19 The `.debug_loclists.dwo` section contains the location lists referenced by  
20 [DW\\_AT\\_location](#) attributes in the `.debug_info.dwo` section. This section has a  
21 similar format to the `.debug_loclists` section in a non-split object, but the  
22 section has some small differences as explained in [Section 7.7.3 on page 237](#).

23 In `demo2.dwo` as shown in [Figure F.7 on the following page](#), the debugging  
24 information for `Line::clip` starting at 2\$ describes a local variable `slope` at 7\$  
25 whose location varies based on the PC. [Figure F.8 on page 429](#) presents some  
26 excerpts from the `.debug_info.dwo` section for `demo2.dwo`.



## Appendix F. Split DWARF Object Files (Informative)

*part 1 of 2*

```
1$: DW_TAG_class_type
    DW_AT_name [DW_FORM_strx]: (slot 20) "Line"
    DW_AT_signature [DW_FORM_ref_sig8]: 0x79c7ef0eae7375d1
    DW_AT_declaration: true
    ...
2$: DW_TAG_subprogram
    DW_AT_external: true
    DW_AT_name [DW_FORM_strx]: (slot 19) "clip"
    DW_AT_decl_file: 2
    DW_AT_decl_line: 16
    DW_AT_linkage_name [DW_FORM_strx]: (slot 2) "_ZN4Line4clipERK3Box"
    DW_AT_type: (reference to DIE for bool)
    DW_AT_accessibility: DW_ACCESS_public
    DW_AT_declaration: true
    ...
```

Figure F.7: Split object example: demo2.dwo DWARF .debug\_info.dwo excerpts

## Appendix F. Split DWARF Object Files (Informative)

part 2 of 2

```
3$: DW_TAG_subprogram
    DW_AT_specification: (reference to 2$)
    DW_AT_decl_file: 1
    DW_AT_decl_line: 3
    DW_AT_low_pc [DW_FORM_addrx]: (slot 32)
    DW_AT_high_pc [DW_FORM_data8]: 0x1ec
    DW_AT_frame_base: DW_OP_call_frame_cfa
    DW_AT_object_pointer: (reference to 4$)
4$: DW_TAG_formal_parameter
    DW_AT_name: (indexed string: 0x11): this
    DW_AT_type: (reference to DIE for type const Point* const)
    DW_AT_artificial: 1
    DW_AT_location: 0x0 (location list)
5$: DW_TAG_formal_parameter
    DW_AT_name: b
    DW_AT_decl_file: 1
    DW_AT_decl_line: 3
    DW_AT_type: (reference to DIE for type const Box& const)
    DW_AT_location [DW_FORM_sec_offset]: 0x2a
6$: DW_TAG_lexical_block
    DW_AT_low_pc [DW_FORM_addrx]: (slot 17)
    DW_AT_high_pc: 0x1d5
7$: DW_TAG_variable
    DW_AT_name [DW_FORM_strx]: (slot 28): "slope"
    DW_AT_decl_file: 1
    DW_AT_decl_line: 5
    DW_AT_type: (reference to DIE for type float)
    DW_AT_location [DW_FORM_sec_offset]: 0x49
```

Figure F.7: Split object example: demo2.dwo DWARF .debug\_info.dwo excerpts  
(concluded)

## Appendix F. Split DWARF Object Files (Informative)

1 In Figure F.7 on page 427, the [DW\\_TAG\\_formal\\_parameter](#) entries at 4\$ and 5\$  
 2 refer to the location lists at offset 0x0 and 0x2a, respectively, and the  
 3 [DW\\_TAG\\_variable](#) entry for slope refers to the location list at offset 0x49. Figure  
 4 F.8 shows a representation of the location lists at those offsets in the  
 5 .debug\_loclists.dwo section.

Entry type		Range		Counted Location Description	
offset	(DW_LLE_*)	start	length	length	expression
0x00	<a href="#">start_length</a>	[9]	0x002f	0x01	<a href="#">DW_OP_reg5</a> (rdi)
0x09	<a href="#">start_length</a>	[11]	0x01b9	0x01	<a href="#">DW_OP_reg3</a> (rbx)
0x12	<a href="#">start_length</a>	[29]	0x0003	0x03	<a href="#">DW_OP_breg12</a> (r12): -8; <a href="#">DW_OP_stack_value</a>
0x1d	<a href="#">start_length</a>	[31]	0x0001	0x03	<a href="#">DW_OP_entry_value</a> : ( <a href="#">DW_OP_reg5</a> (rdi)); <a href="#">DW_OP_stack_value</a>
0x29	<a href="#">end_of_list</a>				
0x2a	<a href="#">start_length</a>	[9]	0x002f	0x01	<a href="#">DW_OP_reg4</a> (rsi)
0x33	<a href="#">start_length</a>	[11]	0x01ba	0x03	<a href="#">DW_OP_reg6</a> (rbp)
0x3c	<a href="#">start_length</a>	[30]	0x0003	0x03	<a href="#">DW_OP_entry_value</a> : ( <a href="#">DW_OP_reg4</a> (rsi)); <a href="#">DW_OP_stack_value</a>
0x48	<a href="#">end_of_list</a>				
0x49	<a href="#">start_length</a>	[10]	0x0004	0x01	<a href="#">DW_OP_reg18</a> (xmm1)
0x52	<a href="#">start_length</a>	[11]	0x01bd	0x02	<a href="#">DW_OP_fbreg</a> : -36
0x5c	<a href="#">end_of_list</a>				

Figure F.8: Split object example: demo2.dwo DWARF .debug\_loclists.dwo excerpts

6 In each [DW\\_LLE\\_start\\_length](#) entry, the start field is the index of a slot in the  
 7 .debug\_addr section, relative to the base offset defined by the compilation unit's  
 8 [DW\\_AT\\_addr\\_base](#) attribute. The .debug\_addr slots referenced by these entries  
 9 give the relocated address of a label within the function where the address range  
 10 begins. The following length field gives the length of the address range. The  
 11 location, consisting of its own length and a DWARF expression, is last.

### 1 **F.3 DWARF Package File Example**

2 A DWARF package file (see Section [7.3.5 on page 198](#)) is a collection of split  
 3 DWARF object files. In general, it will be much smaller than the sum of the split  
 4 DWARF object files, because the packaging process removes duplicate type units  
 5 and merges the string tables. Aside from those two optimizations, however, each  
 6 compilation unit and each type unit from a split DWARF object file is copied  
 7 verbatim into the package file.

8 The package file contains the same set of sections as a split DWARF object file,  
 9 plus two additional sections described below.

10 The packaging utility, like a linker, combines sections of the same name by  
 11 concatenation. While a split DWARF object may contain multiple  
 12 `.debug_info.dwo` sections, one for the compilation unit, and one for each type  
 13 unit, a package file contains a single `.debug_info.dwo` section. The combined  
 14 `.debug_info.dwo` section contains each compilation unit and one copy of each  
 15 type unit (discarding any duplicate type signatures).

16 As part of merging the string tables, the packaging utility treats the  
 17 `.debug_str.dwo` and `.debug_str_offsets.dwo` sections specially. Rather than  
 18 combining them by simple concatenation, it instead builds a new string table  
 19 consisting of the unique strings from each input string table. Because all  
 20 references to these strings use form `DW_FORM_strx`, the packaging utility only  
 21 needs to adjust the string offsets in each `.debug_str_offsets.dwo` contribution  
 22 after building the new `.debug_str.dwo` section.

23 Each compilation unit or type unit consists of a set of inter-related contributions  
 24 to each section in the package file. For example, a compilation unit may have  
 25 contributions in `.debug_info.dwo`, `.debug_abbrev.dwo`, `.debug_line.dwo`,  
 26 `.debug_str_offsets.dwo`, and so on. In order to maintain the ability for a  
 27 consumer to follow references between these sections, the package file contains  
 28 two additional sections: a compilation unit (CU) index, and a type unit (TU)  
 29 index. These indexes allow a consumer to look up a compilation unit (by its  
 30 compilation unit ID) or a type unit (by its type unit signature), and locate each  
 31 contribution that belongs to that unit.

32 For example, consider a package file, `demo.dwp`, formed by combining `demo1.dwo`  
 33 and `demo2.dwo` from the previous example (see Appendix [F.2 on page 417](#)). For  
 34 an executable file named "demo" (or "demo.exe"), a debugger would typically  
 35 expect to find `demo.dwp` in the same directory as the executable file. The resulting  
 36 package file would contain the sections shown in Figure [F.9 on the following](#)  
 37 [page](#), with contributions from each input file as shown.

## Appendix F. Split DWARF Object Files (Informative)

Section	Source of section contributions
<code>.debug_abbrev.dwo</code>	<code>.debug_abbrev.dwo</code> from <code>demo1.dwo</code> <code>.debug_abbrev.dwo</code> from <code>demo2.dwo</code>
<code>.debug_info.dwo</code> (for the compilation units and type units)	compilation unit from <code>demo1.dwo</code> compilation unit from <code>demo2.dwo</code> type unit for class <code>Box</code> from <code>demo1.dwo</code> type unit for class <code>Point</code> from <code>demo1.dwo</code> type unit for class <code>Line</code> from <code>demo2.dwo</code>
<code>.debug_rnglists.dwo</code>	<code>.debug_rnglists.dwo</code> from <code>demo1.dwo</code> <code>.debug_rnglists.dwo</code> from <code>demo2.dwo</code>
<code>.debug_loclists.dwo</code>	<code>.debug_loclists.dwo</code> from <code>demo1.dwo</code> <code>.debug_loclists.dwo</code> from <code>demo2.dwo</code>
<code>.debug_line.dwo</code>	<code>.debug_line.dwo</code> from <code>demo1.dwo</code> <code>.debug_line.dwo</code> from <code>demo2.dwo</code>
<code>.debug_str_offsets.dwo</code>	<code>.debug_str_offsets.dwo</code> from <code>demo1.dwo</code> , adjusted <code>.debug_str_offsets.dwo</code> from <code>demo2.dwo</code> , adjusted
<code>.debug_str.dwo</code>	merged string table generated by package utility
<code>.debug_cu_index</code>	CU index generated by package utility
<code>.debug_tu_index</code>	TU index generated by package utility

Figure F.9: Sections and contributions in example package file `demo.dwp`

1 The `.debug_abbrev.dwo`, `.debug_rnglists.dwo`, `.debug_loclists.dwo` and  
 2 `.debug_line.dwo` sections are copied over from the two `.dwo` files as individual  
 3 contributions to the corresponding sections in the `.dwp` file. The offset of each  
 4 contribution within the combined section and the size of each contribution is  
 5 recorded as part of the CU and TU index sections.

6 The `.debug_info.dwo` sections corresponding to each compilation unit are copied  
 7 as individual contributions to the combined `.debug_info.dwo` section, and one  
 8 copy of each type unit is also copied. The type units for class `Box` and class `Point`,  
 9 for example, are contained in both `demo1.dwo` and `demo2.dwo`, but only one  
 10 instance of each is copied into the package file.

## Appendix F. Split DWARF Object Files (Informative)

1 The `.debug_str.dwo` sections from each file are merged to form a new string  
2 table with no duplicates, requiring the adjustment of all references to those  
3 strings. The `.debug_str_offsets.dwo` sections from the `.dwo` files are copied as  
4 individual contributions, but the string table offset in each slot of those  
5 contributions is adjusted to point to the correct offset in the merged string table.

6 The `.debug_cu_index` and `.debug_tu_index` sections provide a directory to these  
7 contributions. Figure F.10 following shows an example CU index section  
8 containing the two compilation units from `demo1.dwo` and `demo2.dwo`. The CU  
9 index shows that for the compilation unit from `demo1.dwo`, with compilation unit  
10 ID `0x044e413b8a2d1b8f`, its contribution to the `.debug_info.dwo` section begins  
11 at offset 0, and is 325 bytes long. For the compilation unit from `demo2.dwo`, with  
12 compilation unit ID `0xb5f0ecf455e7e97e`, its contribution to the  
13 `.debug_info.dwo` section begins at offset 325, and is 673 bytes long.

14 Likewise, we can find the contributions to the related sections. In Figure F.8 on  
15 page 429, we see that the `DW_TAG_variable` DIE at 7\$ has a reference to a  
16 location list at offset `0x49` (decimal 73). Because this is part of the compilation  
17 unit for `demo2.dwo`, with unit signature `0xb5f0ecf455e7e97e`, we see that its  
18 contribution to `.debug_loclists.dwo` begins at offset 84, so the location list from  
19 Figure F.8 on page 429 can be found in `demo.dwp` at offset 157 ( $84 + 73$ ) in the  
20 combined `.debug_loclists.dwo` section.

21 Figure F.11 following shows an example TU index section containing the three  
22 type units for classes `Box`, `Point`, and `Line`. Each type unit contains contributions  
23 from `.debug_info.dwo`, `.debug_abbrev.dwo`, `.debug_line.dwo` and  
24 `.debug_str_offsets.dwo`. In this example, the type units for classes `Box` and  
25 `Point` come from `demo1.dwo`, and share the abbreviations table, line number  
26 table, and string offsets table with the compilation unit from `demo1.dwo`.  
27 Likewise, the type unit for class `Line` shares tables from `demo2.dwo`.

28 The sharing of these tables between compilation units and type units is typical  
29 for some implementations, but is not required by the DWARF standard.

## Appendix F. Split DWARF Object Files (Informative)

Section header							
Version:							5
Number of columns:							6
Number of used entries:							2
Number of slots:							16

Offset table							
slot	signature	info	abbrev	loc	line	str_off	rng
14	0xb5f0ecf455e7e97e	325	452	84	52	72	350
15	0x044e413b8a2d1b8f	0	0	0	0	0	0

Size table							
slot		info	abbrev	loc	line	str_off	rng
14		673	593	93	52	120	34
15		325	452	84	52	72	15

Figure F.10: Example CU index section

## Appendix F. Split DWARF Object Files (Informative)

Section header					
Version:					5
Number of columns:					4
Number of used entries:					3
Number of slots:					32

Offset table					
slot	signature	info	abbrev	line	str_off
11	0x2f33248f03ff18ab	1321	0	0	0
17	0x79c7ef0eae7375d1	1488	452	52	72
27	0xe97a3917c5a6529b	998	0	0	0

Size table					
slot		info	abbrev	line	str_off
11		167	452	52	72
17		217	593	52	120
27		323	452	52	72

Figure F.11: Example TU index section



## Appendix F. Split DWARF Object Files (Informative)

*(empty page)*

# Appendix G

## DWARF Section Version Numbers (Informative)

Most DWARF sections have a version number in the section header. This version number is not tied to the DWARF standard revision numbers, but instead is incremented when incompatible changes to that section are made. The DWARF standard that a producer is following is not explicitly encoded in the file. Version numbers in the section headers are represented as two-byte unsigned integers.

Table [G.1 on the next page](#) shows what version numbers are in use for each section. In that table:

- “V2” means DWARF Version 2, published July 27, 1993.
- “V3” means DWARF Version 3, published December 20, 2005.
- “V4” means DWARF Version 4, published June 10, 2010.
- “V5” means DWARF Version 5, published February 13, 2017.
- “V6” means DWARF Version 6<sup>1</sup>, published *<to be determined>*.

There are sections with no version number encoded in them; they are only accessed via the `.debug_info` sections and so an incompatible change in those sections’ format would be represented by a change in the `.debug_info` section version number.

---

<sup>1</sup>Higher numbers are reserved for future use.

## Appendix G. Section Version Numbers (Informative)

Table G.1: Section version numbers

Section Name	V2	V3	V4	V5	V6
.debug_abbrev	*	*	*	*	*
.debug_addr	-	-	-	5	5
.debug_aranges	2	2	2	2	2
.debug_frame <sup>2</sup>	1	3	4	4	4
.debug_info	2	3	4	5	5
.debug_line	2	3	4	5	6
.debug_line_str	-	-	-	*	*
.debug_loc	*	*	*	-	-
.debug_loclists	-	-	-	5	5
.debug_macinfo	*	*	*	-	-
.debug_macro	-	-	-	5	5
.debug_names	-	-	-	5	6
.debug_pubnames	2	2	2	-	-
.debug_pubtypes	-	2	2	-	-
.debug_ranges	-	*	*	-	-
.debug_rnglists	-	-	-	5	5
.debug_str	*	*	*	*	*
.debug_str_offsets	-	-	-	5	5
.debug_sup	-	-	-	5	5
.debug_types	-	-	4	-	-
<i>(split object sections)</i>					
.debug_abbrev.dwo	-	-	-	*	*
.debug_info.dwo	-	-	-	5	5
.debug_line.dwo	-	-	-	5	5
.debug_loclists.dwo	-	-	-	5	5
.debug_macro.dwo	-	-	-	5	5
.debug_rnglists.dwo	-	-	-	5	5
.debug_str.dwo	-	-	-	*	*
.debug_str_offsets.dwo	-	-	-	5	5

*Continued on next page*

<sup>2</sup>For the `.debug_frame` section, version 2 is unused.

## Appendix G. Section Version Numbers (Informative)

<b>Section Name</b>	<b>V2</b>	<b>V3</b>	<b>V4</b>	<b>V5</b>	<b>V6</b>
<i>(package file sections)</i>					
<code>.debug_cu_index</code>	-	-	-	5	6
<code>.debug_tu_index</code>	-	-	-	5	6

1 Notes:

- 2 • “\*” means that a version number is not applicable (the section does not  
3 include a header or the section’s header does not include a version).
- 4 • “-” means that the section was not defined in that version of the DWARF  
5 standard.
- 6 • The version numbers for corresponding `.debug_<kind>` and  
7 `.debug_<kind>.dwo` sections are the same.

Appendix G. Section Version Numbers (Informative)

*(empty page)*

# Appendix H

## GNU Free Documentation License

Version 1.3, 3 November 2008

Copyright ©2000, 2001, 2002, 2007, 2008 Free Software Foundation, Inc.

<http://fsf.org/>

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

### **PREAMBLE**

The purpose of this License is to make a manual, textbook, or other functional and useful document “free” in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of “copyleft,” which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

## H.1 APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The “Document”, below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as “you.” You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A “Modified Version” of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A “Secondary Section” is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document’s overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The “Invariant Sections” are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The “Cover Texts” are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A “Transparent” copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or

## Appendix H. GNU Free Documentation License

discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not “Transparent” is called “Opaque.”

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The “Title Page” means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, “Title Page” means the text near the most prominent appearance of the work’s title, preceding the beginning of the body of the text.

The “publisher” means any person or entity that distributes copies of the Document to the public.

A section “Entitled XYZ” means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as “Acknowledgements,” “Dedications,” “Endorsements,” or “History.”) To “Preserve the Title” of such a section when you modify the Document means that it remains a section “Entitled XYZ” according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

## H.2 VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License.



You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section [H.3](#).

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

### **H.3 COPYING IN QUANTITY**

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

## H.4 MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections [H.2](#) and [H.3](#) above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.
- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section Entitled "History," Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations

## Appendix H. GNU Free Documentation License

given in the Document for previous versions it was based on. These may be placed in the “History” section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.

- K. For any section Entitled “Acknowledgements” or “Dedications”, Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- M. Delete any section Entitled “Endorsements”. Such a section may not be included in the Modified Version.
- N. Do not retitle any existing section to be Entitled “Endorsements” or to conflict in title with any Invariant Section.
- O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version’s license notice. These titles must be distinct from any other section titles.

You may add a section Entitled “Endorsements,” provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

## H.5 COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section [H.5](#) above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled “History” in the various original documents, forming one section Entitled “History;” likewise combine any sections Entitled “Acknowledgements,” and any sections Entitled “Dedications.” You must delete all sections Entitled “Endorsements.”

## H.6 COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

## H.7 AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an “aggregate” if the copyright resulting from the compilation

is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section [H.3](#) is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

## H.8 TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section [H.4](#). Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section [H.4](#)) to Preserve its Title (section [H.1](#)) will typically require changing the actual title.

## H.9 TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, or distribute it is void, and will automatically terminate your rights under this License.

However, if you cease all violation of this License, then your license from a particular copyright holder is reinstated (a) provisionally, unless and until the copyright holder explicitly and finally terminates your license, and (b) permanently, if the copyright holder fails to notify you of the violation by some reasonable means prior to 60 days after the cessation.

## Appendix H. GNU Free Documentation License

Moreover, your license from a particular copyright holder is reinstated permanently if the copyright holder notifies you of the violation by some reasonable means, this is the first time you have received notice of violation of this License (for any work) from that copyright holder, and you cure the violation prior to 30 days after your receipt of the notice.

Termination of your rights under this section does not terminate the licenses of parties who have received copies or rights from you under this License. If your rights have been terminated and not permanently reinstated, receipt of a copy of some or all of the same material does not give you any rights to use it.

### **H.10 FUTURE REVISIONS OF THIS LICENSE**

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <http://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License “or any later version” applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation. If the Document specifies that a proxy can decide which future versions of this License can be used, that proxy’s public statement of acceptance of a version permanently authorizes you to choose that version for the Document.

### **H.11 RELICENSING**

“Massive Multiauthor Collaboration Site” (or “MMC Site”) means any World Wide Web server that publishes copyrightable works and also provides prominent facilities for anybody to edit those works. A public wiki that anybody can edit is an example of such a server. A “Massive Multiauthor Collaboration” (or “MMC”) contained in the site means any set of copyrightable works thus published on the MMC site.

“CC-BY-SA” means the Creative Commons Attribution-Share Alike 3.0 license published by Creative Commons Corporation, a not-for-profit corporation with a

## Appendix H. GNU Free Documentation License

principal place of business in San Francisco, California, as well as future copyleft versions of that license published by that same organization.

“Incorporate” means to publish or republish a Document, in whole or in part, as part of another Document.

An MMC is “eligible for relicensing” if it is licensed under this License, and if all works that were first published under this License somewhere other than this MMC, and subsequently incorporated in whole or in part into the MMC, (1) had no cover texts or invariant sections, and (2) were thus incorporated prior to November 1, 2008.

The operator of an MMC Site may republish an MMC contained in the site under CC-BY-SA on the same site at any time before August 1, 2009, provided the MMC is eligible for relicensing.

### **ADDENDUM: How to use this License for your documents**

To use this License in a document you have written, include a copy of the License in the document and put the following copyright and license notices just after the title page:

Copyright (C) YEAR YOUR NAME.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.3 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts.

A copy of the license is included in the section entitled “GNU Free Documentation License.”

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the “with... Texts.” line with this:

with the Invariant Sections being LIST THEIR TITLES, with the Front-Cover Texts being LIST, and with the Back-Cover Texts being LIST.

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

## Appendix H. GNU Free Documentation License

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.



Appendix H. GNU Free Documentation License

*(empty page)*

# Index

- &-qualified non-static member
  - function, 21
- &&-qualified non-static member
  - function, 22
- <caf>, *see* code alignment factor
- <daf>, *see* data alignment factor
- ... parameters, *see* unspecified parameters entry
- .data, 387
- .debug\_abbrev.dwo, 8, 196, 198, 199, 202, 290, 291, 413, 423, 430–432, 437
- .debug\_abbrev, 147, 193, 195, 206, 210–213, 286, 287, 290, 291, 298, 299, 387, 397, 414, 420, 422, 437
  - example, 298
- .debug\_addr, 8, 28, 45, 54, 55, 69, 194, 195, 223, 224, 253, 286, 288–290, 292, 414, 415, 420–423, 429, 437
- .debug\_aranges, 154, 192, 194, 195, 206, 246, 286, 287, 290, 291, 387, 392, 414, 415, 420–422, 437
- .debug\_cu\_index, 8, 199, 200, 431, 432, 438
- .debug\_frame, 181–183, 192, 195, 206, 286, 290, 343, 414, 437
- .debug\_info.dwo, 8, 15, 196–199, 202, 205, 290–293, 413, 415, 423, 424, 426–428, 430–432, 437
- .debug\_info, 8, 9, 13, 15, 24, 28, 31, 36, 41, 69, 141, 142, 144, 150, 151, 154, 155, 182, 192–195, 197, 205–211, 213, 217, 227–230, 246, 286–292, 298, 299, 387, 388, 390, 392–398, 409, 410, 413, 414, 420, 421, 423, 436, 437
  - example, 298
- .debug\_line.dwo, 8, 72, 165, 174, 177, 196, 198, 199, 202, 225, 290, 292, 293, 413, 423, 424, 430–432, 437
- .debug\_line\_str, 8, 155, 160, 163–165, 195, 207, 229, 286, 289, 290, 414, 437
- .debug\_line, 65, 155, 160, 164, 165, 174, 177, 192, 194, 195, 206, 207, 225, 286–290, 292, 380, 381, 387, 392, 397, 413, 414, 420–422, 437
- .debug\_loclists.dwo, 8, 43, 196, 198, 199, 202, 255, 290, 292, 413, 423, 426, 429, 431, 432, 437
- .debug\_loclists, 9, 10, 43, 46, 69, 192, 194, 208, 226, 255, 286, 288, 292, 413, 426, 437
- .debug\_loc (pre-Version 5), 10, 437
- .debug\_macinfo (pre-Version 5), 8, 199, 437

## Index

- .debug\_macro.dwo, 8, 71, 174, 196, 198, 199, 202, 226, 290, 292, 293, 413, 423, 437
- .debug\_macro, 8, 65, 172, 174, 178, 194, 204, 226, 286–289, 380, 381, 413, 437
- .debug\_names, 9, 10, 142, 149, 150, 192, 194, 195, 206, 207, 286, 287, 290, 291, 414, 415, 420–422, 437
- .debug\_pubnames (pre-Version 5), 9, 10, 142, 437
- .debug\_pubtypes (pre-Version 5), 9, 10, 142, 437
- .debug\_ranges (pre-Version 5), 10, 437
- .debug\_rnglists.dwo, 53, 196, 198, 199, 202, 254, 290, 292, 423, 426, 431, 437
- .debug\_rnglists, 9, 10, 53, 55, 56, 69, 192, 194, 208, 227, 254, 286, 288–290, 292, 437
- .debug\_str.dwo, 8, 152, 165, 196–199, 290, 293, 413, 423, 426, 430–432, 437
- .debug\_str\_offsets.dwo, 8, 68, 165, 197–199, 202, 205, 253, 290, 291, 293, 413, 414, 423, 426, 430–432, 437
- .debug\_str\_offsets, 8, 68, 72, 175, 194, 196, 205, 207, 229, 230, 253, 285–288, 290, 291, 415, 437
- .debug\_str, 152, 164, 165, 175, 178, 194–196, 204, 205, 207, 229, 230, 286–288, 290, 291, 379, 381, 414, 415, 420–422, 437
- .debug\_sup, 203, 437
- .debug\_tu\_index, 8, 199, 200, 431, 432, 438
- .debug\_types (Version 4), 8, 437
- .dwo file extension, 413
- .dwp file extension, 198
- .text, 387, 394, 396
- segment\_selector\_size (deprecated), 154, 160, 182, 247, 254, 255
- 32-bit DWARF format, 36, 160, 204, 210–212, 224–230, 246, 251, 256
- 64-bit DWARF format, 36, 160, 204, 210–212, 225–230, 246, 251, 256, 257
- abbrev\_table\_size, 150, 153, 246
- abbreviations table, 208
  - dynamic forms in, 217
  - example, 298
- abstract instance, 20, 410
  - entry, 85
  - example, 348, 351
  - nested, 90
  - root, 85
  - tree, 85
- abstract origin attribute, 87–89, 219
- accelerated access, 141
  - by address, 154
  - by name, 142
- Access declaration, 17
- access declaration entry, 122
- accessibility attribute, 17, 47, 121, 122, 124, 219, 240
- Accessibility of base or inherited class, 17
- activation of call frame, 178, 189
- Ada, 1, 17, 47, 64, 110, 112, 113, 138, 305, 318, 320, 324, 325, 328, 329
- address, *see also* address class
  - dereference operator, 30, 31
  - implicit push for member operator, 135

## Index

- implicit push of base, 32
  - uplevel, *see* static link attribute
- address class, 23, 45, 49, 56, 90, 218, 220–223, 230, 231, 243
- address class attribute, 82, 114, 219
- address index, 45, 54
- address of call instruction, 18
- address of called routine, 19
- address of called routine, which may be clobbered, 19
- address of the value pointed to by an argument, 18
- address register
  - in line number machine, 157
- address size, *see also* `address_size`, *see* size of an address
- address space
  - multiple, 31
- address table, 17
- address table base
  - encoding, 222
- address table base attribute, 69
- `address_range`, 182, 183, 189
- `address_size`, 9, 154, 160, 182, 210, 211, 246, 247, 254, 255, 339
- `addrptr`, *see also* `addrptr` class
- `addrptr` class, 23, 69, 222–224, 230, 231
- adjusted opcode, 169
- alias declaration, *see* imported declaration entry
- alignment
  - non-default, 17
- alignment attribute, 59, 222
- all calls summary attribute, 80, 222
- all source calls summary attribute, 81, 222
- all tail and normal calls are described, 18
- all tail calls are described, 18
- all tail calls summary attribute, 80, 222
- all tail, normal and inlined calls are described, 18
- allocated attribute, 138, 220
- allocation status of types, 17
- anonymous structure, 318
- anonymous union, 99, 122
- ANSI-defined language names, 63, 241
- argument value passed, 19
- ARM instruction set architecture, 155
- array
  - assumed-rank, 139, 313
  - declaration of type, 115
  - descriptor for, 304
  - element ordering, 115
  - element type, 116
- Array bound THREADS scale factor, 22
- array coarray, *see* coarray
- array element stride (of array type), 18
- array row/column ordering, 21
- array type entry, 115
  - examples, 304
- artificial attribute, 48, 219
- artificial name or description, 20
- ASCII (Fortran string kind), 132
- ASCII (Fortran string kind), 109
- ASCII character, 108
- Assembly, 64
- associated attribute, 138, 220
- associated compilation unit, 223
- association status of types, 17
- assumed-rank array, *see* array, assumed-rank
- atomic qualified type entry, 113
- attribute duplication, 17
- attribute encodings, 218
- attribute ordering, 17

## Index

- attribute value classes, 23
- attributes, 15
  - list of, 17
- augmentation, 150, 182
- augmentation sequence, 150
- augmentation string, 182
- augmentation\_size, 150
- auto return type, 81, 112, 126
  
- base address, 52
  - of location list, 44
  - of range list, 54
- base address of scope, 21, 52
- base address selection entry
  - in range list, 252, 253
- base type bit location, 19
- base type bit size, 18
- base type entry, 106
- base types attribute, 68, 219
- basic block, 156, 157, 159, 167, 169
- basic\_block, 157, 159, 167, 169
- beginning of a data member, 123
- beginning of an object, 123, 323
- Bernstein hash function, 262
- Bias added to an encoded value, 17
- bias attribute, 223
- big-endian encoding, *see* endianness
  - attribute
- binary scale attribute, 109, 221
- binary scale factor for fixed-point type, 17
- bit fields, 321, 323
- bit offset attribute (Version 3), 218
- bit size attribute, 106, 107, 124, 136, 218
- bit stride attribute, 116, 130, 135, 219
- BLISS, 64
- block, *see also* block class, 25, 82, 98, 128, 189
- block class, 23, 219, 220, 224, 230
  
- bounded location description, 43, 45, 46
- bounded range, 53, 55, 56
- bucket\_count, 150, 151
- byte order, 102, 198, 203, 225, 232, 323
- byte size attribute, 106, 124, 136, 218
- byte stride attribute, 116, 130, 135, 220
  
- C, 1, 12, 21, 32, 49, 64, 68, 79–81, 95, 98–100, 107, 109, 112–116, 118, 119, 130, 131, 136, 172, 177, 262, 313, 318, 323, 365, 371, 379, 393, 394, 409
- C++, 1, 9, 11, 12, 17, 21, 22, 32, 47, 48, 51, 58, 64, 74–77, 81, 83–85, 87, 89, 91, 95, 97, 99, 100, 102, 109, 112, 113, 115, 118–122, 124–126, 130, 131, 135–137, 149, 172, 177, 318, 323, 331, 353, 356, 359–363, 383, 386, 389, 391–394, 398, 409
- C++ for OpenCL, 64
- C++11, 125, 131
- C-interoperable, 313
- C#, 64
- call column attribute, 94, 221
  - of call site entry, 94
- call data location attribute, 94, 222
- call data value attribute, 94, 222
- call file attribute, 94, 221
  - of call site entry, 94
- call is a tail call, 18
- call line attribute, 94, 221
  - of call site entry, 94
- call origin attribute, 93, 222
- call parameter attribute, 95, 222
- call PC attribute, 222
- call pc attribute, 93
- call return PC attribute, 222
- call return pc attribute, 93

## Index

- call site
  - address of called routine, [19](#)
  - address of called routine, which may be clobbered, [19](#)
  - address of the call instruction, [18](#)
  - address of the value pointed to by an argument, [18](#)
  - argument value passed, [19](#)
  - parameter entry, [18](#)
  - return address, [18](#)
  - subprogram called, [18](#)
  - summary
    - all tail and normal calls are described, [18](#)
    - all tail calls are described, [18](#)
    - all tail, normal and inlined calls are described, [18](#)
  - tail call, [18](#)
  - value pointed to by an argument, [18](#)
- call site entry, [92](#)
- call site parameter entry, [94](#)
- call site return pc attribute, [93](#)
- call site summary information, [80](#)
- call tail call attribute, [93](#), [222](#)
- call target attribute, [93](#), [222](#)
- call target clobbered attribute, [93](#), [222](#)
- call type attribute, [93](#)
- call value attribute, [94](#), [222](#)
- Calling convention
  - for subprograms, [19](#)
  - for types, [19](#)
- calling convention attribute, [219](#)
  - for subprogram, [79](#)
  - for types, [120](#)
- calling convention codes
  - for subroutines, [79](#)
  - for types, [120](#)
- catch block, [97](#)
- catch block entry, [97](#)
- Child determination encodings, [217](#)
- CIE\_id, [182](#), [206](#), [207](#)
- CIE\_pointer, [182](#), [183](#), [206](#), [207](#)
- class of attribute value
  - address, *see* address class
  - addrptr, *see* addrptr class
  - block, *see* block class
  - constant, *see* constant class
  - exprloc, *see* exprloc class
  - flag, *see* flag class
  - lineptr, *see* lineptr class
  - loclist, *see* loclist class
  - loclistsptr, *see* loclistsptr class
  - macptr, *see* macptr class
  - reference, *see* reference class
  - rnglist, *see* rnglist class
  - rnglistsptr, *see* rnglistsptr class
  - string, *see* string class
  - stroffsetsptr, *see* stroffsetsptr class
- class type entry, [118](#)
- class variable entry, [124](#)
- coarray, [117](#)
  - example, [310–312](#)
- COBOL, [1](#), [3](#), [12](#), [64](#)
- code address or range of addresses, [21](#)
- code alignment factor, [182](#)
- code\_alignment\_factor, [182](#), [184](#)
- codimension, *see* coarray
- coindex, *see* coarray
- column, [157](#)
- column position of call site of non-inlined call, [18](#)
- column position of inlined subroutine call, [18](#)
- column position of source declaration, [19](#)
- COMDAT, [11](#), [389](#), [396–398](#)
- common, [82](#)
- common block, *see* Fortran common block, [99](#)
- common block entry, [102](#)

## Index

- common block reference attribute, [82](#),  
[102](#)
- common block usage, [19](#)
- common blocks, [51](#)
- common information entry, [181](#)
- common reference attribute, [219](#)
- comp\_unit\_count, [150](#)
- compilation directory, [19](#)
- compilation directory attribute, [66](#),  
[219](#)
- compilation unit, [61](#)
  - see also* type unit, [71](#)
  - full, [62](#)
  - partial, [62](#)
  - skeleton, [69](#)
- compilation unit ID, [430](#), [432](#)
- compilation unit set, [199](#)
- compilation unit uses UTF-8 strings,  
[22](#)
- compile-time constant function, [19](#)
- compile-time constant object, [19](#)
- compiler identification, [21](#)
- concrete instance
  - example, [348](#), [351](#), [352](#)
  - nested, [90](#)
- concrete out-of-line instance, [410](#)
- condition entry, [129](#)
- const qualified type entry, [113](#)
- constant, *see also* constant class, [57](#)
- constant (data) entry, [99](#)
- constant class, [23](#), [56](#), [59](#), [98](#), [101](#), [107](#),  
[195](#), [218–224](#), [230](#), [231](#), [245](#)
- constant expression attribute, [86](#), [102](#),  
[221](#)
- constant object, [19](#)
- constant value attribute, [59](#), [101](#), [130](#),  
[219](#)
- constexpr, [84](#), [87](#)
- containing type (of pointer) attribute,  
[135](#)
- containing type attribute, [219](#)
- containing type of pointer to member  
type, [19](#)
- contiguous range of code addresses,  
[20](#)
- conventional compilation unit, *see*  
full compilation unit, partial  
compilation unit, type unit
- conventional type unit, [71](#)
- count attribute, [114](#), [134](#), [220](#)
  - default, [134](#)
- counted location description, [44](#)
- Crystal, [64](#)
- D, [64](#), [113](#)
- data (indirect) location attribute, [137](#)
- data alignment factor, [183](#)
- data bit offset, [323](#)
- data bit offset attribute, [107](#), [123](#), [221](#)
- data bit size, [323](#)
- data location attribute, [220](#)
- data member, *see* member entry  
(data)
- data member attribute, [220](#)
- data member bit location, [19](#)
- data member bit size, [18](#)
- data member location, [19](#)
- data member location attribute, [121](#),  
[123](#)
- data object entries, [99](#)
- data object location, [21](#)
- data object or data type size, [18](#)
- data\_alignment\_factor, [183](#),  
[185–187](#)
- debug\_abbrev\_offset, [9](#), [202](#), [206](#),  
[210–212](#), [287](#), [291](#)
- debug\_info\_offset, [154](#), [194](#), [206](#),  
[246](#)
- debug\_line\_offset, [173](#), [174](#)
- debug\_line\_offset\_flag, [173](#), [380](#),  
[381](#)
- debugging information entry, [15](#)

## Index

- ownership relation, 25, 390
- debugging information entry
  - relationship, 22
- decimal scale attribute, 109–112, 221
- decimal scale factor, 19
- decimal sign attribute, 110, 111
- decimal sign representation, 19
- DECL, 263–283
- declaration attribute, 49, 50, 73, 99, 100, 119, 220
- declaration column attribute, 50, 51, 220
- declaration coordinates, 15, 50, 87, 263
  - in concrete instance, 87
- declaration file attribute, 50, 220
- declaration line attribute, 50, 220
- DEFAULT (Fortran string kind), 132
- default location description, 44
- default value attribute, 101, 219
- default value of parameter, 19
- default\_is\_stmt, 159, 161, 339
- defaulted attribute, 19, 126, 222
- deleted attribute, 126, 222
- Deletion of member function, 19
- denominator of rational scale factor, 22
- derived type (C++), *see* inheritance entry
- description attribute, 57, 221
- descriptor
  - array, 304
- DIE, *see* debugging information entry
- digit count attribute, 109–111, 221
- digit count for packed decimal or numeric string type, 20
- directories, 162, 165, 339
- directories\_count, 162, 339
- directory\_format\_count, 162, 339
- directory\_format\_table, 162, 339
- discontiguous address ranges, *see* non-contiguous address ranges
- discriminant (entry), 127
- discriminant attribute, 127, 218
- discriminant list attribute, 128, 220, 245
- discriminant of variant part, 20
- discriminant value, 20
- discriminant value attribute, 128, 218
- discriminated union, *see* variant entry
- discriminator, 156, 158, 159, 167, 169, 172
- divisor of rational scale factor, 222
- DJB hash function, 151, 262
- duplication elimination, *see* DWARF duplicate elimination
- DW\_ACCESS\_private, 47, 240, 401, 403, 404
- DW\_ACCESS\_protected, 47, 240
- DW\_ACCESS\_public, 47, 240, 424, 427
- DW\_ADDR\_none, 49, 243
- DW\_AT\_abstract\_origin, 17, 87, 88, 89, 143, 219, 263, 350, 354, 356, 357, 368
- DW\_AT\_accessibility, 17, 47, 121, 122, 124, 219, 240, 259, 264, 266, 268–273, 276–283, 401, 403, 404, 424, 427
- DW\_AT\_addr\_base, 17, 28, 45, 54, 69, 70, 71, 196, 222, 224, 254, 267, 274, 276, 286, 288, 292, 415, 416, 420, 422, 423, 429
- DW\_AT\_address\_class, 17, 49, 82, 114, 219, 259, 268, 274, 275, 278, 280, 283
- DW\_AT\_alignment, 17, 59, 222, 259, 264–271, 273–282
- DW\_AT\_allocated, 17, 117, 137, 138, 138, 220, 259, 264–266,



## Index

- 268–270, 275–277, 279–282,  
309, 317
- DW\_AT\_artificial, 15, 17, 48, 90, 125,  
219, 259, 263, 325, 330, 336,  
338, 402, 408, 425, 428
- DW\_AT\_associated, 17, 117, 137, 138,  
138, 220, 259, 264–266,  
268–270, 275–277, 279–282,  
308, 317
- DW\_AT\_base\_types, 17, 68, 70, 71,  
219, 267, 274, 416
- DW\_AT\_bias, 17, 107, 223, 265, 325
- DW\_AT\_binary\_scale, 17, 109, 109,  
221, 259, 265
- DW\_AT\_bit\_offset (deprecated), 218
- DW\_AT\_bit\_size, 18, 58, 58, 106, 107,  
116, 119, 124, 130, 132–134,  
136, 218, 259, 264–266, 269,  
270, 272, 274–277, 279, 282,  
321–323, 325
- DW\_AT\_bit\_stride, 18, 58, 116, 130,  
130, 135, 219, 259, 264, 269,  
270, 279, 322
- DW\_AT\_byte\_size, 10, 18, 58, 58, 106,  
116, 119, 124, 130, 132–134,  
136, 218, 259, 264–266, 269,  
270, 272, 274–277, 279, 282,  
299, 309, 325, 357, 358, 366,  
370, 397, 399–403, 405, 408,  
425
- DW\_AT\_byte\_stride, 18, 58, 116, 130,  
130, 135, 220, 259, 269, 270,  
279, 307, 314
- DW\_AT\_call\_all\_calls, 18, 80, 80, 81,  
222
- DW\_AT\_call\_all\_source\_calls, 18, 80,  
81, 81, 222
- DW\_AT\_call\_all\_tail\_calls, 18, 80, 80,  
81, 222
- DW\_AT\_call\_column, 18, 86, 94, 221,  
265, 271
- DW\_AT\_call\_data\_location, 18, 94,  
94, 222, 266
- DW\_AT\_call\_data\_value, 18, 94, 94,  
222, 266, 378
- DW\_AT\_call\_file, 18, 86, 94, 221, 265,  
271
- DW\_AT\_call\_line, 18, 86, 94, 221, 265,  
271
- DW\_AT\_call\_origin, 18, 93, 93, 222,  
265, 374, 378
- DW\_AT\_call\_parameter, 18, 95, 222,  
266, 374, 375, 378
- DW\_AT\_call\_pc, 18, 81, 92, 93, 222,  
265
- DW\_AT\_call\_return\_pc, 18, 81, 92,  
93, 222, 265, 374, 375, 378
- DW\_AT\_call\_tail\_call, 18, 93, 222, 265
- DW\_AT\_call\_target, 19, 93, 93, 222,  
265, 374, 375
- DW\_AT\_call\_target\_clobbered, 19,  
93, 93, 222, 265
- DW\_AT\_call\_value, 19, 94, 94, 222,  
266, 374, 375, 378
- DW\_AT\_calling\_convention, 19, 79,  
120, 219, 244, 266, 277, 278,  
282
- DW\_AT\_common\_reference, 19, 82,  
219, 267, 396, 397
- DW\_AT\_comp\_dir, 19, 63, 66, 66,  
69–71, 219, 267, 274, 276, 299,  
415, 416, 420, 422, 424
- DW\_AT\_const\_expr, 19, 86, 87, 102,  
221, 259, 271, 282, 357
- DW\_AT\_const\_value, 19, 41, 57, 59,  
86, 87, 101, 130, 219, 259,  
268–270, 281, 282, 350, 357,  
359, 370, 395
- DW\_AT\_containing\_type, 19, 135,  
219, 259, 275, 338
- DW\_AT\_count, 19, 114, 134, 220, 259,  
270, 276, 279

## Index

- DW\_AT\_data\_bit\_offset, 19, 57, 107, 123, 123, 221, 259, 265, 272, 321–323, 325
- DW\_AT\_data\_location, 19, 117, 137, 137, 138, 220, 259, 264–266, 268–270, 275–277, 279–282, 305–310, 314, 317, 319, 370
- DW\_AT\_data\_member\_location, 19, 32, 57, 121, 123, 123, 220, 259, 271, 272, 309, 320, 327, 329, 330, 366, 399–401, 403–405
- DW\_AT\_decimal\_scale, 19, 109, 110, 111, 112, 221, 259, 265
- DW\_AT\_decimal\_sign, 19, 110, 111, 221, 239, 259, 265
- DW\_AT\_decl\_column, 19, 50, 51, 220, 261, 263
- DW\_AT\_decl\_file, 19, 50, 50, 72, 196, 220, 261, 263, 399, 401, 402, 408, 413, 424, 425, 427, 428
- DW\_AT\_decl\_line, 19, 50, 50, 220, 261, 263, 399, 401, 402, 408, 424, 425, 427, 428
- DW\_AT\_declaration, 19, 49, 50, 73, 99, 100, 119, 143, 220, 260, 261, 264, 266–270, 272, 273, 275–283, 335, 336, 402, 408, 424, 427
- DW\_AT\_default\_value, 19, 59, 101, 219, 259, 270, 280, 281, 383, 384
- DW\_AT\_defaulted, 9, 19, 126, 126, 222, 246, 278
- DW\_AT\_deleted, 9, 19, 126, 222, 278
- DW\_AT\_description, 15, 20, 57, 221, 261, 263
- DW\_AT\_digit\_count, 20, 109, 110, 111, 112, 221, 259, 265
- DW\_AT\_discr, 20, 127, 218, 259, 283, 327–330
- DW\_AT\_discr\_list, 20, 128, 128, 220, 245, 259, 283
- DW\_AT\_discr\_value, 20, 128, 128, 218, 259, 283, 327, 329, 330
- DW\_AT\_dwo\_name, 20, 69, 69, 70, 211, 222, 274, 276, 291, 415, 416, 420, 422
- DW\_AT\_elemental, 20, 80, 221, 278
- DW\_AT\_encoding, 20, 106, 107, 220, 238, 259, 265, 299, 325, 358, 370, 397, 399, 400, 402, 403, 408, 425
- DW\_AT\_endianity, 20, 102, 102, 106, 221, 240, 259, 265, 268, 270, 282
- DW\_AT\_entry\_pc, 20, 56, 56, 68, 71, 74, 81, 85, 86, 95–97, 143, 220, 266, 267, 271, 272, 274, 278, 281, 283, 416
- DW\_AT\_enum\_class, 20, 130, 130, 221, 259, 269, 359
- DW\_AT\_explicit, 20, 124, 221, 259, 278
- DW\_AT\_export\_symbols, 20, 74, 75, 118, 222, 266, 273, 277, 282, 318, 332, 334
- DW\_AT\_extension, 20, 74, 221, 273, 333
- DW\_AT\_external, 20, 78, 99, 100, 153, 154, 220, 268, 278, 282, 402, 408, 424, 427
- DW\_AT\_frame\_base, 20, 29, 33, 82, 83, 220, 268, 278, 301, 354–356, 425, 428
- DW\_AT\_friend, 20, 122, 220, 259, 260, 270
- DW\_AT\_hi\_user, 191, 223
- DW\_AT\_high\_pc, 11, 20, 52, 53, 63, 70, 71, 74, 81, 85, 86, 95–97, 143, 218, 266, 267, 271, 272, 274, 276, 278, 281, 283, 299, 332, 333, 350, 351, 354–356,

## Index

- 375, 408, 415, 416, 423, 425, 428
- DW\_AT\_identifer\_case, 20, 67, 71, 220, 243, 267, 274, 395, 416
- DW\_AT\_import, 20, 75, 76, 76, 77, 78, 204, 219, 271, 333, 391, 393, 396, 397
- DW\_AT\_inline, 20, 84, 84, 85, 85, 219, 244, 278, 348, 349, 352, 355, 357, 368
- DW\_AT\_is\_optional, 20, 101, 219, 259, 270
- DW\_AT\_language (deprecated), 218
- DW\_AT\_language\_name, 21, 63, 71, 72, 101, 116, 223, 241, 267, 274, 281, 299, 393, 395, 399, 401, 416, 424
- DW\_AT\_language\_version, 21, 65, 71, 72, 101, 223, 267, 274, 281, 416
- DW\_AT\_linkage\_name, 21, 51, 58, 78, 102, 143, 221, 267, 268, 278, 282, 424, 427
- DW\_AT\_lo\_user, 191, 223
- DW\_AT\_location, 21, 36, 41, 51, 51, 59, 85, 94, 96, 100, 102, 143, 218, 259, 266–268, 270, 282, 283, 286, 309, 317, 320, 322, 325, 332, 333, 336, 350, 351, 354–356, 366, 368, 370, 374, 375, 378, 395, 406, 425, 426, 428
- DW\_AT\_loclists\_base, 21, 43, 69, 222, 256, 416
- DW\_AT\_low\_pc, 11, 21, 52, 52, 53, 63, 69–71, 74, 81, 85, 86, 95–97, 143, 218, 266–268, 271, 272, 274, 276, 278, 281, 283, 299, 332, 333, 350, 351, 354–356, 375, 408, 415, 416, 423, 425, 428
- DW\_AT\_lower\_bound, 21, 134, 219, 241, 259, 270, 279, 308, 309, 311, 314, 320, 322, 325, 395
- DW\_AT\_macro\_info, 65, 66
- DW\_AT\_macro\_info (deprecated), 220
- DW\_AT\_macros, 21, 65, 65, 71, 222, 267, 274, 286, 292, 416
- DW\_AT\_main\_subprogram, 11, 21, 68, 68, 71, 78, 221, 267, 274, 278, 416
- DW\_AT\_mutable, 21, 123, 221, 259, 272
- DW\_AT\_name, 21, 51, 51, 57, 59, 63, 67, 71, 74, 76, 78, 88, 95, 96, 99, 102, 103, 106, 112, 113, 115, 117, 118, 120, 122, 129–133, 135–137, 143, 164, 218, 258–260, 263–283, 299, 309, 311, 312, 317, 318, 320–322, 325, 327, 329, 330, 332–336, 338, 349, 351, 352, 355, 357–361, 363, 364, 366, 368, 370, 375, 383, 384, 393–397, 399–406, 408, 416, 424, 425, 427, 428
- DW\_AT\_namelist\_item, 21, 103, 220, 273
- DW\_AT\_noreturn, 9, 21, 80, 222, 278
- DW\_AT\_object\_pointer, 21, 125, 125, 221, 261, 278, 335, 336, 425, 428
- DW\_AT\_ordering, 21, 115, 218, 244, 259, 264, 311, 312
- DW\_AT\_picture\_string, 21, 111, 221, 259, 265
- DW\_AT\_priority, 21, 74, 220, 272
- DW\_AT\_producer, 21, 67, 71, 219, 267, 274, 299, 416, 424
- DW\_AT\_prototyped, 21, 80, 131, 219, 259, 278, 280

## Index

- DW\_AT\_pure, 21, 80, 221, 278
- DW\_AT\_ranges, 21, 52, 52, 53, 63, 70, 71, 74, 81, 85, 86, 95–97, 143, 221, 266, 267, 271, 272, 274, 276, 278, 281, 283, 286, 292, 415, 416, 423, 426
- DW\_AT\_rank, 9, 21, 117, 139, 139, 221, 259, 264, 313, 314
- DW\_AT\_recursive, 21, 80, 80, 221, 278
- DW\_AT\_reference, 21, 125, 131, 132, 222, 259, 278, 280
- DW\_AT\_return\_addr, 21, 82, 85, 219, 268, 271, 278
- DW\_AT\_rnglists\_base, 22, 53, 69, 70, 222, 255, 267, 274, 276, 286, 416
- DW\_AT\_rvalue\_reference, 22, 125, 131, 132, 222, 259, 278, 280, 338
- DW\_AT\_scale\_divisor, 22, 110, 110, 222, 259, 265
- DW\_AT\_scale\_multiplier, 22, 110, 110, 222, 259, 265
- DW\_AT\_segment (deprecated), 220
- DW\_AT\_sibling, 22, 25, 50, 218, 263
- DW\_AT\_signature, 22, 50, 119, 221, 266, 269, 271, 277, 280, 282, 408, 424, 427
- DW\_AT\_small, 22, 110, 110, 221, 259, 265
- DW\_AT\_specification, 22, 50, 50, 75, 85, 100, 119, 125, 126, 143, 220, 260, 264, 266, 269, 272, 277, 278, 282, 333, 408, 425, 428
- DW\_AT\_start\_scope, 22, 85, 97, 219, 264, 266, 268, 269, 271–273, 276, 277, 279–282
- DW\_AT\_static\_link, 22, 83, 83, 220, 268, 279, 350, 354, 355
- DW\_AT\_stmt\_list, 22, 65, 70–72, 202, 218, 267, 274, 276, 281, 286, 292, 299, 415, 416, 420–424
- DW\_AT\_str\_offsets, 22, 68, 70–72, 222, 253, 267, 274, 276, 281, 286, 287, 415, 416
- DW\_AT\_str\_offsets\_base (deprecated), 22
- DW\_AT\_string\_length, 10, 22, 132, 132, 133, 219, 259, 277, 370
- DW\_AT\_string\_length\_bit\_size, 22, 58, 132, 221, 259, 277
- DW\_AT\_string\_length\_byte\_size, 22, 58, 132, 221, 259, 277, 370
- DW\_AT\_tensor, 22, 115, 116, 223
- DW\_AT\_threads\_scaled, 22, 134, 221, 259, 270, 279
- DW\_AT\_trampoline, 22, 90, 221, 271, 279
- DW\_AT\_type, 22, 47, 47, 59, 81, 83, 93, 93, 95, 96, 100, 102, 113, 115–117, 121, 122, 127, 129–131, 132, 133–137, 220, 259, 260, 264–276, 279–283, 299, 308, 309, 311, 312, 314, 317, 320–322, 325, 327, 329, 330, 332, 333, 335, 336, 338, 349, 352, 355, 357–364, 366, 368, 370, 375, 383, 384, 393–396, 399–406, 408, 424, 425, 427, 428
- DW\_AT\_upper\_bound, 22, 134, 134, 219, 259, 270, 279, 308, 309, 311, 312, 314, 320, 322, 325, 357, 395
- DW\_AT\_use\_location, 22, 135, 135, 136, 220, 259, 275
- DW\_AT\_use\_UTF8, 22, 68, 70–72, 220, 229, 259, 267, 274, 276, 281, 416
- DW\_AT\_variable\_parameter, 22, 100, 220, 259, 270
- DW\_AT\_virtuality, 23, 48, 121, 124,

## Index

- 220, 241, 259, 271, 279
- DW\_AT\_visibility, 23, 47, 219, 240, 259, 264, 266–270, 272, 273, 275–277, 279–283
- DW\_AT\_vtable\_elem\_location, 23, 125, 220, 259, 279
- DW\_ATE\_address, 108, 238
- DW\_ATE\_ASCII, 108, 109, 132, 239, 370
- DW\_ATE\_boolean, 108, 238, 425
- DW\_ATE\_complex\_float, 108, 110, 238
- DW\_ATE\_complex\_signed, 108, 112, 239
- DW\_ATE\_complex\_unsigned, 108, 112, 239
- DW\_ATE\_decimal\_float, 108, 110, 239
- DW\_ATE\_edited, 108, 111, 239
- DW\_ATE\_float, 108, 110, 238
- DW\_ATE\_hi\_user, 191, 239
- DW\_ATE\_imaginary\_float, 108, 110, 239
- DW\_ATE\_imaginary\_signed, 108, 112, 239
- DW\_ATE\_imaginary\_unsigned, 108, 112, 239
- DW\_ATE\_lo\_user, 191, 239
- DW\_ATE\_numeric\_string, 108, 109, 110, 112, 239
- DW\_ATE\_packed\_decimal, 108, 109, 110, 112, 239
- DW\_ATE\_signed, 107, 108, 239, 325, 397, 399, 400, 402, 403, 408
- DW\_ATE\_signed\_bitint, 108, 109, 239
- DW\_ATE\_signed\_char, 108, 239
- DW\_ATE\_signed\_fixed, 108, 109, 239
- DW\_ATE\_UCS, 108, 109, 132, 239, 370
- DW\_ATE\_unsigned, 108, 239, 325
- DW\_ATE\_unsigned\_bitint, 108, 109, 239
- DW\_ATE\_unsigned\_char, 108, 239, 299
- DW\_ATE\_unsigned\_fixed, 108, 109, 239
- DW\_ATE\_UTF, 108, 109, 239, 358
- DW\_CC\_hi\_user, 191, 244
- DW\_CC\_lo\_user, 191, 244
- DW\_CC\_nocall, 79, 79, 244
- DW\_CC\_normal, 79, 79, 120, 244
- DW\_CC\_pass\_by\_reference, 120, 244
- DW\_CC\_pass\_by\_value, 120, 244
- DW\_CC\_program, 79, 79, 244
- DW\_CFA\_advance\_loc, 184, 184, 185, 189, 251, 346
- DW\_CFA\_advance\_loc1, 185, 185, 251
- DW\_CFA\_advance\_loc2, 185, 185, 251
- DW\_CFA\_advance\_loc4, 185, 185, 251
- DW\_CFA\_def\_cfa, 185, 185, 251, 345
- DW\_CFA\_def\_cfa\_expression, 184, 186, 186, 251
- DW\_CFA\_def\_cfa\_offset, 185, 185, 186, 251, 346
- DW\_CFA\_def\_cfa\_offset\_sf, 186, 186, 252
- DW\_CFA\_def\_cfa\_register, 185, 185, 251, 346
- DW\_CFA\_def\_cfa\_sf, 185, 185, 252
- DW\_CFA\_expression, 184, 187, 187, 251
- DW\_CFA\_hi\_user, 191, 252
- DW\_CFA\_lo\_user, 191, 252
- DW\_CFA\_nop, 183, 184, 188, 188, 251, 345, 346
- DW\_CFA\_offset, 186, 186, 251, 346
- DW\_CFA\_offset\_extended, 186, 186, 187, 251

## Index

DW\_CFA\_offset\_extended\_sf, 187, 187, 252  
DW\_CFA\_register, 187, 187, 251, 345  
DW\_CFA\_remember\_state, 188, 188, 251  
DW\_CFA\_restore, 188, 188, 251, 346  
DW\_CFA\_restore\_extended, 188, 188, 251  
DW\_CFA\_restore\_state, 188, 188, 251  
DW\_CFA\_same\_value, 186, 186, 251, 345  
DW\_CFA\_set\_loc, 184, 184, 189, 251  
DW\_CFA\_undefined, 186, 186, 189, 251, 345  
DW\_CFA\_val\_expression, 184, 188, 188, 252  
DW\_CFA\_val\_offset, 187, 187, 252  
DW\_CFA\_val\_offset\_sf, 187, 187, 252  
DW\_CHILDREN\_no, 216, 217, 299  
DW\_CHILDREN\_yes, 216, 217, 299  
DW\_DEFAULTED\_in\_class, 126, 246  
DW\_DEFAULTED\_no, 126, 246  
DW\_DEFAULTED\_out\_of\_class, 126, 246  
DW\_DS\_leading\_overpunch, 111, 240  
DW\_DS\_leading\_separate, 111, 240  
DW\_DS\_trailing\_overpunch, 111, 240  
DW\_DS\_trailing\_separate, 111, 240  
DW\_DS\_unsigned, 111, 240  
DW\_DSC\_label, 128, 245  
DW\_DSC\_range, 128, 245  
DW\_END\_big, 102, 240  
DW\_END\_default, 102, 240  
DW\_END\_hi\_user, 191, 240  
DW\_END\_little, 102, 240  
DW\_END\_lo\_user, 191, 240  
DW\_FORM\_addr, 45, 55, 193, 223, 230, 299  
DW\_FORM\_addrx, 9, 69, 195, 196, 223, 224, 231, 286, 288, 292, 414, 415, 422, 423, 425, 428  
DW\_FORM\_addrx1, 9, 69, 195, 196, 224, 224, 231, 288, 292, 414, 415, 422, 423  
DW\_FORM\_addrx2, 9, 69, 195, 196, 224, 224, 231, 288, 292, 414, 415, 422, 423  
DW\_FORM\_addrx3, 9, 69, 195, 196, 224, 224, 231, 288, 292, 414, 415, 422, 423  
DW\_FORM\_addrx4, 9, 69, 195, 196, 224, 224, 231, 288, 292, 414, 415, 422, 423  
DW\_FORM\_addrx\_offset, 217, 217  
DW\_FORM\_block, 165, 166, 174, 224, 230, 258, 407  
DW\_FORM\_block1, 166, 174, 224, 230  
DW\_FORM\_block2, 166, 174, 224, 230  
DW\_FORM\_block4, 166, 174, 224, 230  
DW\_FORM\_data, 225, 225  
DW\_FORM\_data1, 165, 166, 174, 224, 225, 230, 299  
DW\_FORM\_data16, 9, 166, 174, 224, 225, 231  
DW\_FORM\_data2, 165, 166, 174, 224, 225, 230  
DW\_FORM\_data4, 11, 165, 166, 174, 224, 225, 230, 384  
DW\_FORM\_data8, 11, 165, 166, 174, 224, 225, 230, 245, 425, 428  
DW\_FORM\_data<n>, 107  
DW\_FORM\_exprloc, 231  
DW\_FORM\_exprval, 225, 258  
DW\_FORM\_flag, 166, 174, 225, 230, 258, 407  
DW\_FORM\_flag\_present, 225, 231, 246  
DW\_FORM\_implicit\_const, 9, 217, 217, 225, 231  
DW\_FORM\_indirect, 217, 217, 231, 299



## Index

- DW\_FORM\_line\_strp, 9, 164–166, 174, 207, 229, 229, 231, 286, 289, 293, 413
- DW\_FORM\_locdesc, 226
- DW\_FORM\_loclistx, 9, 46, 69, 226, 231, 256, 288, 292
- DW\_FORM\_ref1, 195, 227, 231, 390
- DW\_FORM\_ref2, 36, 195, 227, 231, 390
- DW\_FORM\_ref4, 36, 195, 227, 231, 299, 390
- DW\_FORM\_ref8, 195, 227, 231, 390
- DW\_FORM\_ref<n>, 227, 393
- DW\_FORM\_ref\_addr, 12, 36, 194, 197, 207, 227, 228, 231, 286, 287, 299, 388, 390–392, 410
- DW\_FORM\_ref\_sig8, 144, 151, 212, 228, 231, 399, 424, 427
- DW\_FORM\_ref\_sup4, 9, 204, 228, 231
- DW\_FORM\_ref\_sup8, 9, 204, 228, 231
- DW\_FORM\_ref\_udata, 195, 227, 231, 390
- DW\_FORM\_rnglistx, 9, 55, 69, 227, 231, 255, 288, 292
- DW\_FORM\_sdata, 166, 174, 224, 225, 230, 258, 383, 400, 403–405, 407
- DW\_FORM\_sec\_offset, 11, 46, 56, 166, 174, 194, 197, 207, 223, 224–227, 230, 231, 255, 256, 287, 288, 292, 299, 428
- DW\_FORM\_string, 164–166, 174, 195, 228, 230, 258, 299, 339, 383, 384, 400, 403–405, 407, 424, 425
- DW\_FORM\_strp, 164–166, 174, 194, 195, 197, 207, 229, 229, 230, 253, 286, 287, 291, 339, 414, 426
- DW\_FORM\_strp8, 164, 166, 174, 194, 195, 197, 229, 229, 231, 287, 291, 414
- DW\_FORM\_strp\_sup, 9, 164, 166, 174, 204, 207, 229, 229, 231
- DW\_FORM\_strp\_sup8, 164, 166, 174, 204, 229, 229, 232
- DW\_FORM\_strx, 10, 68, 72, 165, 166, 174, 195–197, 199, 229, 229, 231, 286, 287, 291, 413–415, 423–425, 427, 428, 430
- DW\_FORM\_strx1, 10, 68, 72, 165, 166, 174, 195–197, 199, 229, 229, 231, 287, 291, 413–415, 423
- DW\_FORM\_strx2, 10, 68, 72, 165, 166, 174, 195–197, 199, 229, 229, 231, 287, 291, 413–415, 423
- DW\_FORM\_strx3, 10, 68, 72, 165, 166, 174, 195–197, 199, 229, 229, 231, 287, 291, 413–415, 423
- DW\_FORM\_strx4, 10, 68, 72, 165, 166, 174, 195–197, 199, 229, 229, 231, 287, 291, 413–415, 423
- DW\_FORM\_udata, 165, 166, 174, 224, 225, 231, 339
- DW\_ID\_case\_insensitive, 67, 68, 243, 395
- DW\_ID\_case\_sensitive, 67, 67, 243
- DW\_ID\_down\_case, 67, 67, 243
- DW\_ID\_up\_case, 67, 67, 243
- DW\_IDX\_compile\_unit, 153, 245
- DW\_IDX\_die\_offset, 153, 245
- DW\_IDX\_external, 153, 154, 245, 246
- DW\_IDX\_hi\_user, 191, 245
- DW\_IDX\_lo\_user, 191, 245
- DW\_IDX\_parent, 153, 245
- DW\_IDX\_type\_hash, 153, 245
- DW\_IDX\_type\_unit, 153, 245
- DW\_INL\_declared\_inlined, 85, 244, 349, 352, 355, 368
- DW\_INL\_declared\_not\_inlined, 85, 244
- DW\_INL\_inlined, 84, 85, 244, 357

## Index

DW\_INL\_not\_inlined, [85](#), [85](#), [244](#)  
DW\_LLE\_base\_address, [46](#), [238](#)  
DW\_LLE\_base\_addressx, [45](#), [69](#), [238](#),  
[288](#)  
DW\_LLE\_default\_location, [46](#), [238](#)  
DW\_LLE\_end\_of\_list, [45](#), [46](#), [238](#),  
[368](#), [429](#)  
DW\_LLE\_hi\_user, [191](#), [238](#)  
DW\_LLE\_include\_loclist, [46](#), [238](#)  
DW\_LLE\_include\_loclistx, [46](#), [238](#)  
DW\_LLE\_lo\_user, [191](#), [238](#)  
DW\_LLE\_offset\_pair, [45](#), [45](#), [46](#), [238](#)  
DW\_LLE\_start\_end, [46](#), [238](#), [368](#)  
DW\_LLE\_start\_length, [46](#), [238](#), [429](#)  
DW\_LLE\_startx\_endx, [45](#), [69](#), [238](#),  
[288](#)  
DW\_LLE\_startx\_length, [45](#), [69](#), [238](#),  
[288](#)  
DW\_LNAME\_Ada, [64](#), [242](#)  
DW\_LNAME\_Assembly, [64](#), [243](#)  
DW\_LNAME\_BLISS, [64](#), [242](#)  
DW\_LNAME\_C, [64](#), [242](#), [299](#)  
DW\_LNAME\_C\_plus\_plus, [64](#), [242](#),  
[393](#), [399](#), [401](#), [424](#)  
DW\_LNAME\_C\_sharp, [64](#), [243](#)  
DW\_LNAME\_Cobol, [64](#), [242](#)  
DW\_LNAME\_CPP\_for\_OpenCL, [64](#),  
[243](#)  
DW\_LNAME\_Crystal, [64](#), [242](#)  
DW\_LNAME\_D, [64](#), [242](#)  
DW\_LNAME\_Dylan, [64](#), [242](#)  
DW\_LNAME\_Fortran, [64](#), [242](#), [395](#)  
DW\_LNAME\_GLSL, [64](#), [243](#)  
DW\_LNAME\_GLSL\_ES, [64](#), [243](#)  
DW\_LNAME\_Go, [64](#), [242](#)  
DW\_LNAME\_Haskell, [64](#), [242](#)  
DW\_LNAME\_hi\_user, [191](#), [243](#)  
DW\_LNAME\_HLSL, [64](#), [243](#)  
DW\_LNAME\_Java, [64](#), [242](#)  
DW\_LNAME\_Julia, [64](#), [242](#)  
DW\_LNAME\_Kotlin, [64](#), [242](#)  
DW\_LNAME\_lo\_user, [191](#), [243](#)  
DW\_LNAME\_Modula2, [64](#), [242](#)  
DW\_LNAME\_Modula3, [64](#), [242](#)  
DW\_LNAME\_Mojo, [64](#), [243](#)  
DW\_LNAME\_ObjC, [64](#), [242](#)  
DW\_LNAME\_ObjC\_plus\_plus, [64](#),  
[242](#)  
DW\_LNAME\_OCaml, [64](#), [242](#)  
DW\_LNAME\_OpenCL\_C, [65](#), [242](#)  
DW\_LNAME\_OpenCL\_CPP, [65](#), [243](#)  
DW\_LNAME\_Pascal, [65](#), [242](#)  
DW\_LNAME\_PLI, [65](#), [242](#)  
DW\_LNAME\_Python, [65](#), [242](#)  
DW\_LNAME\_RenderScript, [65](#), [242](#)  
DW\_LNAME\_Ruby, [65](#), [243](#)  
DW\_LNAME\_Rust, [65](#), [242](#)  
DW\_LNAME\_Swift, [65](#), [242](#)  
DW\_LNAME\_SYCL, [65](#), [243](#)  
DW\_LNAME\_UPC, [65](#), [242](#)  
DW\_LNAME\_Zig, [65](#), [242](#)  
DW\_LNCT\_directory\_index, [165](#),  
[248](#), [339](#)  
DW\_LNCT\_hi\_user, [166](#), [191](#), [248](#)  
DW\_LNCT\_lo\_user, [166](#), [191](#), [248](#)  
DW\_LNCT\_MD5, [166](#), [166](#), [248](#)  
DW\_LNCT\_path, [164](#), [166](#), [248](#), [339](#)  
DW\_LNCT\_size, [166](#), [166](#), [248](#), [339](#)  
DW\_LNCT\_source, [166](#), [166](#), [248](#), [339](#)  
DW\_LNCT\_timestamp, [165](#), [165](#), [248](#),  
[339](#)  
DW\_LNCT\_URL, [166](#), [166](#), [248](#), [339](#)  
DW\_LNE\_define\_file (deprecated),  
[248](#)  
DW\_LNE\_end\_sequence, [171](#), [171](#),  
[172](#), [248](#), [342](#)  
DW\_LNE\_hi\_user, [191](#), [248](#)  
DW\_LNE\_lo\_user, [191](#), [248](#)  
DW\_LNE\_padding, [172](#), [172](#), [248](#)  
DW\_LNE\_set\_address, [172](#), [172](#), [194](#),  
[248](#), [414](#)



## Index

- DW\_LNE\_set\_discriminator, 172, 172, 248
- DW\_LNS\_advance\_line, 169, 169, 247
- DW\_LNS\_advance\_pc, 169, 169, 170, 247, 342
- DW\_LNS\_const\_add\_pc, 170, 170, 247
- DW\_LNS\_copy, 169, 169, 247
- DW\_LNS\_extended\_op, 159, 171, 171, 247
- DW\_LNS\_fixed\_advance\_pc, 159, 170, 170, 247, 342
- DW\_LNS\_hi\_user, 191
- DW\_LNS\_lo\_user, 191
- DW\_LNS\_negate\_stmt, 161, 169, 169, 247
- DW\_LNS\_set\_basic\_block, 169, 169, 247
- DW\_LNS\_set\_column, 169, 169, 247
- DW\_LNS\_set\_epilogue\_begin, 171, 171, 247
- DW\_LNS\_set\_file, 169, 169, 247
- DW\_LNS\_set\_isa, 171, 171, 247
- DW\_LNS\_set\_prologue\_end, 170, 170, 247
- DW\_MACRO\_define, 175, 175, 176, 250, 380, 381
- DW\_MACRO\_define\_strp, 175, 175, 176, 178, 194, 197, 204, 250, 286, 288, 381
- DW\_MACRO\_define\_strx, 175, 175, 176, 197, 250, 286, 288, 293
- DW\_MACRO\_define\_sup (deprecated), 250
- DW\_MACRO\_define\_sup4, 173, 175, 175, 176, 204, 250
- DW\_MACRO\_define\_sup8, 173, 175, 175, 176, 204, 250
- DW\_MACRO\_end\_file, 174, 177, 177, 250, 380, 381
- DW\_MACRO\_hi\_user, 174, 191, 250
- DW\_MACRO\_import, 178, 178, 250, 286, 289, 381
- DW\_MACRO\_import\_sup (deprecated), 250
- DW\_MACRO\_import\_sup4, 173, 178, 178, 204, 250
- DW\_MACRO\_import\_sup8, 173, 178, 178, 204, 250
- DW\_MACRO\_lo\_user, 174, 191, 250
- DW\_MACRO\_start\_file, 174, 177, 177, 196, 250, 286, 288, 293, 380–382
- DW\_MACRO\_undef, 175, 175, 176, 250, 380, 381
- DW\_MACRO\_undef\_strp, 175, 175, 176, 178, 194, 197, 204, 250, 286, 288
- DW\_MACRO\_undef\_strx, 175, 175, 176, 197, 250, 286, 288, 293
- DW\_MACRO\_undef\_sup (deprecated), 250
- DW\_MACRO\_undef\_sup4, 173, 175, 175, 176, 204, 250
- DW\_MACRO\_undef\_sup8, 173, 175, 175, 176, 204, 250
- DW\_OP\_abs, 33, 33, 234
- DW\_OP\_addr, 27, 27, 143, 193, 234, 301
- DW\_OP\_addrx, 28, 28, 69, 184, 195, 196, 236, 286, 288, 292, 414
- DW\_OP\_and, 33, 33, 234, 308, 309
- DW\_OP\_bit\_piece, 42, 42, 43, 236, 304
- DW\_OP\_bra, 36, 36, 235
- DW\_OP\_breg0, 29, 235, 373, 378
- DW\_OP\_breg1, 29, 235, 302, 303
- DW\_OP\_breg31, 29, 235, 301
- DW\_OP\_breg<n>, 29, 82
- DW\_OP\_bregx, 29, 29, 40, 236, 301
- DW\_OP\_call2, 36, 36, 51, 184, 236
- DW\_OP\_call4, 36, 36, 51, 184, 236

## Index

DW\_OP\_call\_frame\_cfa, 33, 33, 184, 236, 425, 428  
DW\_OP\_call\_ref, 36, 36, 51, 57, 184, 197, 207, 236, 261, 286, 287  
DW\_OP\_const1s, 27, 234  
DW\_OP\_const1u, 27, 234  
DW\_OP\_const2s, 27, 234  
DW\_OP\_const2u, 27, 234  
DW\_OP\_const4s, 27, 234  
DW\_OP\_const4u, 27, 234  
DW\_OP\_const8s, 27, 234  
DW\_OP\_const8u, 27, 234  
DW\_OP\_const<n><x>, 27, 33  
DW\_OP\_const<n>s, 27  
DW\_OP\_const<n>u, 27  
DW\_OP\_const\_type, 27, 28, 28, 184, 236  
DW\_OP\_consts, 27, 27, 234  
DW\_OP\_constu, 27, 27, 234  
DW\_OP\_constx, 28, 28, 69, 184, 195, 196, 236, 286, 288, 292  
DW\_OP\_convert, 37, 37, 184, 237  
DW\_OP\_deref, 30, 30, 234, 301, 303, 308, 309, 314, 374, 375  
DW\_OP\_deref\_size, 30, 30, 31, 236, 378  
DW\_OP\_deref\_type, 31, 31, 184, 237  
DW\_OP\_div, 33, 34, 34, 234  
DW\_OP\_drop, 30, 30, 234, 300  
DW\_OP\_dup, 29, 29, 234, 300  
DW\_OP\_entry\_value, 37, 37, 38, 236, 303, 373–375, 378, 429  
DW\_OP\_eq, 35, 235  
DW\_OP\_fpreg, 29, 29, 236, 301, 302, 425, 429  
DW\_OP\_form\_tls\_address, 32, 32, 33, 143, 236  
DW\_OP\_ge, 35, 235  
DW\_OP\_gt, 35, 235  
DW\_OP\_hi\_user, 38, 191, 237  
DW\_OP\_implicit\_pointer, 41, 41, 236, 365, 366, 368  
DW\_OP\_implicit\_value, 40, 40, 236, 365, 368  
DW\_OP\_le, 35, 235  
DW\_OP\_lit0, 27, 235  
DW\_OP\_lit1, 27, 235, 302, 308, 368, 374  
DW\_OP\_lit31, 27, 235  
DW\_OP\_lit<n>, 27, 34, 308, 309, 314, 320  
DW\_OP\_lo\_user, 38, 191, 237  
DW\_OP\_lt, 35, 235  
DW\_OP\_minus, 33, 34, 34, 234  
DW\_OP\_mod, 34, 34, 235  
DW\_OP\_mul, 33, 34, 34, 235, 314, 373, 375  
DW\_OP\_ne, 35, 235  
DW\_OP\_neg, 33, 34, 34, 235  
DW\_OP\_nop, 37, 37, 236  
DW\_OP\_not, 34, 34, 235  
DW\_OP\_or, 34, 34, 235  
DW\_OP\_over, 30, 30, 234, 300  
DW\_OP\_pick, 30, 30, 234, 300  
DW\_OP\_piece, 42, 42, 43, 236, 302, 304, 366, 368  
DW\_OP\_plus, 33, 34, 34, 235, 302, 308, 309, 314, 320  
DW\_OP\_plus\_uconst, 34, 34, 235, 301, 303, 329  
DW\_OP\_push\_object\_address, 32, 32, 37, 94, 123, 138, 184, 236, 307–310, 314  
DW\_OP\_reg0, 40, 235, 302, 304, 373–375, 378  
DW\_OP\_reg1, 40, 235, 303, 304, 373–375  
DW\_OP\_reg31, 40, 235  
DW\_OP\_reg<n>, 40, 82  
DW\_OP\_regval\_bits, 28, 29, 29, 237

## Index

- DW\_OP\_regval\_type, 28, 29, 29, 184, 236
- DW\_OP\_regx, 40, 40, 236, 301
- DW\_OP\_reinterpret, 37, 37, 184, 237
- DW\_OP\_rot, 30, 30, 234, 300
- DW\_OP\_shl, 35, 35, 235
- DW\_OP\_shr, 35, 35, 235
- DW\_OP\_shra, 35, 35, 235
- DW\_OP\_skip, 36, 36, 235
- DW\_OP\_stack\_value, 40, 40, 236, 302, 303, 365, 366, 368, 373, 375, 429
- DW\_OP\_swap, 30, 30, 234, 300
- DW\_OP\_user\_extended, 38, 38, 237
- DW\_OP\_xderef, 31, 31, 234
- DW\_OP\_xderef\_size, 31, 31, 32, 236
- DW\_OP\_xderef\_type, 32, 32, 237
- DW\_OP\_xor, 35, 35, 235
- DW\_ORD\_col\_major, 116, 245, 311, 312
- DW\_ORD\_row\_major, 116, 245
- DW\_RLE\_base\_address, 55, 252
- DW\_RLE\_base\_addressx, 54, 69, 252, 286, 289
- DW\_RLE\_end\_of\_list, 54, 55, 56, 252
- DW\_RLE\_hi\_user, 191, 252
- DW\_RLE\_include\_rnglist, 56, 252
- DW\_RLE\_include\_rnglistx, 55, 252
- DW\_RLE\_lo\_user, 191, 252
- DW\_RLE\_offset\_pair, 54, 55, 55, 252
- DW\_RLE\_start\_end, 55, 252
- DW\_RLE\_start\_length, 56, 252
- DW\_RLE\_startx\_endx, 55, 69, 252, 286, 289
- DW\_RLE\_startx\_length, 55, 69, 252, 286, 289
- DW\_SECT\_ABBREV, 202
- DW\_SECT\_INFO, 202
- DW\_SECT\_LINE, 202
- DW\_SECT\_LOCLISTS, 202
- DW\_SECT\_MACRO, 202
- DW\_SECT\_RNGLISTS, 202
- DW\_SECT\_STR\_OFFSETS, 202
- DW\_TAG\_access\_declaration, 16, 122, 215, 264
- DW\_TAG\_array\_type, 16, 115, 214, 264, 308, 309, 311, 312, 314, 320, 322, 357, 395
- DW\_TAG\_atomic\_type, 16, 113, 216, 264
- DW\_TAG\_base\_type, 16, 28, 29, 31, 32, 37, 106, 107, 114, 132, 215, 265, 299, 321, 325, 332, 335, 358, 370, 393, 397, 399, 400, 402, 403, 408, 425
- DW\_TAG\_call\_site, 9, 16, 80, 81, 92, 92, 216, 265, 374, 375, 378
- DW\_TAG\_call\_site\_parameter, 16, 94, 94, 216, 266, 374, 375, 378
- DW\_TAG\_catch\_block, 16, 97, 215, 266
- DW\_TAG\_class\_type, 16, 118, 127, 214, 266, 335, 338, 401, 403, 408, 424, 427
- DW\_TAG\_coarray\_type, 9, 16, 117, 216, 266, 311, 312
- DW\_TAG\_common\_block, 16, 102, 214, 267, 395
- DW\_TAG\_common\_inclusion, 16, 82, 214, 267, 396, 397
- DW\_TAG\_compile\_unit, 16, 62, 62, 71, 208, 214, 267, 287, 288, 292, 299, 391–394, 396, 397, 406, 413, 423, 424
- DW\_TAG\_condition, 3, 16, 129, 216, 267
- DW\_TAG\_const\_type, 16, 113, 114, 215, 267, 335, 338, 357, 370, 425
- DW\_TAG\_constant, 16, 99, 99, 110, 129, 215, 268, 370, 395
- DW\_TAG\_dwarf\_procedure, 16, 41,

## Index

- 51, 215, 268, 368
- DW\_TAG\_dynamic\_type, 16, 136, 137, 216, 268, 317
- DW\_TAG\_entry\_point, 16, 78, 214, 268
- DW\_TAG\_enumeration\_type, 16, 116, 129, 130, 214, 269, 359
- DW\_TAG\_enumerator, 16, 130, 215, 269, 359
- DW\_TAG\_file\_type, 16, 136, 215, 269
- DW\_TAG\_formal\_parameter, 16, 41, 93, 95, 97, 99, 129, 131, 214, 270, 336, 338, 349, 350, 352, 354–357, 361, 366, 368, 383, 384, 402, 408, 425, 428, 429
- DW\_TAG\_friend, 16, 122, 215, 259, 260, 270
- DW\_TAG\_generic\_subrange, 9, 16, 117, 117, 134, 139, 216, 270, 313, 314
- DW\_TAG\_hi\_user, 191, 216
- DW\_TAG\_immutable\_type, 16, 113, 216, 270
- DW\_TAG\_imported\_declaration, 16, 75, 214, 271, 333
- DW\_TAG\_imported\_module, 16, 75, 76, 215, 271, 333
- DW\_TAG\_imported\_unit, 16, 78, 204, 215, 271, 391, 393, 396, 397, 410
- DW\_TAG\_inheritance, 16, 121, 214, 271
- DW\_TAG\_inlined\_subroutine, 16, 78, 81, 86, 87–90, 92, 143, 214, 271, 349, 350, 354, 356, 357, 368
- DW\_TAG\_interface\_type, 16, 120, 215, 271
- DW\_TAG\_label, 16, 96, 143, 214, 272
- DW\_TAG\_lexical\_block, 16, 95, 214, 272, 373, 375, 408, 428
- DW\_TAG\_lo\_user, 191, 216
- DW\_TAG\_member, 16, 93, 122, 129, 214, 272, 309, 317, 318, 320–322, 325, 327, 329, 330, 360, 361, 363, 366, 393, 399–401, 403–405
- DW\_TAG\_module, 16, 74, 215, 272
- DW\_TAG\_mutable\_type (deprecated), 216
- DW\_TAG\_namelist, 16, 103, 215, 273
- DW\_TAG\_namelist\_item, 16, 103, 215, 273
- DW\_TAG\_namespace, 16, 74, 74, 143, 215, 273, 332–334, 399–405, 408
- DW\_TAG\_packed\_type, 16, 113, 215, 273, 321
- DW\_TAG\_partial\_unit, 16, 62, 62, 208, 215, 274, 287, 288, 292, 391, 392, 395, 396
- DW\_TAG\_pointer\_type, 16, 113, 114, 214, 259, 260, 274, 299, 335, 338, 402, 404, 408, 425
- DW\_TAG\_ptr\_to\_member\_type, 16, 135, 215, 259, 260, 275, 338
- DW\_TAG\_reference\_type, 16, 113, 114, 214, 259, 260, 275, 394, 425
- DW\_TAG\_restrict\_type, 16, 113, 114, 215, 275
- DW\_TAG\_rvalue\_reference\_type, 16, 113, 114, 216, 259, 260, 275
- DW\_TAG\_set\_type, 16, 133, 215, 276
- DW\_TAG\_shared\_type, 16, 113, 114, 216, 276
- DW\_TAG\_skeleton\_unit, 16, 69, 216, 276, 420, 422
- DW\_TAG\_string\_type, 16, 132, 214, 277, 370
- DW\_TAG\_structure\_type, 16, 118, 127, 214, 277, 309, 317, 318, 320–322, 325, 327, 329, 330,

## Index

- 360, 361, 363, 364, 366, 393,  
399, 400, 402, 405
- DW\_TAG\_subprogram, 16, 78, 78, 84,  
88–90, 113, 124, 143, 215, 259,  
278, 279, 332, 333, 335, 336,  
338, 349, 350, 352, 354–357,  
361, 362, 366, 368, 373, 383,  
384, 394, 396, 397, 402, 405,  
406, 408, 424, 425, 427, 428
- DW\_TAG\_subrange\_type, 16, 116,  
117, 129, 133, 139, 215, 241,  
279, 308, 309, 311, 312, 320,  
322, 325, 357, 395
- DW\_TAG\_subroutine\_type, 16, 131,  
214, 280, 338
- DW\_TAG\_template\_alias, 16, 137,  
216, 280, 363, 364
- DW\_TAG\_template\_type\_parameter,  
16, 59, 215, 280, 360, 361, 363,  
364
- DW\_TAG\_template\_value\_parameter,  
16, 59, 215, 281
- DW\_TAG\_thrown\_type, 16, 83, 215,  
281
- DW\_TAG\_try\_block, 16, 97, 215, 281
- DW\_TAG\_type\_unit, 16, 62, 72, 208,  
216, 281, 287, 292, 399, 401,  
413, 424
- DW\_TAG\_typedef, 16, 115, 115, 214,  
281, 299
- DW\_TAG\_union\_type, 16, 118, 127,  
214, 282
- DW\_TAG\_unspecified\_parameters,  
16, 82, 97, 131, 214, 282
- DW\_TAG\_unspecified\_type, 16, 112,  
215, 282
- DW\_TAG\_variable, 16, 41, 88, 93, 99,  
100, 114, 124, 129, 143, 215,  
282, 309, 311, 312, 317, 320,  
322, 325, 332–334, 338, 349,  
350, 352, 354–361, 363, 364,  
366, 368, 370, 375, 394–396,  
406, 428, 429, 432
- DW\_TAG\_variant, 16, 128, 214, 283,  
327, 329, 330
- DW\_TAG\_variant\_part, 16, 127, 215,  
283, 327–330
- DW\_TAG\_volatile\_type, 16, 113, 114,  
215, 283
- DW\_TAG\_with\_stmt, 16, 96, 215, 283
- DW\_UT\_compile, 209, 210
- DW\_UT\_hi\_user, 191, 209
- DW\_UT\_lo\_user, 191, 209
- DW\_UT\_partial, 209, 210
- DW\_UT\_skeleton, 209, 210
- DW\_UT\_split\_compile, 209, 210
- DW\_UT\_split\_type, 209, 211
- DW\_UT\_type, 209, 211
- DW\_VIRTUALLY\_none, 48, 241
- DW\_VIRTUALLY\_pure\_virtual, 48,  
241
- DW\_VIRTUALLY\_virtual, 48, 241
- DW\_VIS\_exported, 48, 241
- DW\_VIS\_local, 48, 241
- DW\_VIS\_qualified, 48, 241
- DWARF compression, 386, 412
- DWARF duplicate elimination, 386,  
*see also* DWARF compression,  
*see also* split DWARF object  
file  
examples, 392–394, 430
- DWARF expression, 27, *see also*  
location description  
arithmetic operations, 33  
control flow operations, 35  
examples, 300  
literal encodings, 27  
logical operations, 33  
operator encoding, 233  
register based addressing, 28  
special operations, 37  
stack operations, 27, 29

## Index

- DWARF package file, [430](#)
- DWARF package files, [198](#)
  - section identifier encodings, [202](#)
- DWARF procedure, [51](#)
- DWARF procedure entry, [51](#)
- DWARF Version 1, [iii](#)
- DWARF Version 2, [iii](#), [2](#), [12](#), [13](#), [161](#), [193](#), [436](#)
- DWARF Version 3, [iii](#), [11](#), [12](#), [161](#), [168](#), [216](#), [218](#), [436](#)
- DWARF Version 4, [1](#), [2](#), [8](#), [10](#), [11](#), [27](#), [161](#), [199](#), [220](#), [248](#), [436](#)
- DWARF Version 5, [1](#), [8](#), [9](#), [43](#), [53](#), [61](#), [65](#), [133](#), [142](#), [154](#), [160](#), [161](#), [163](#), [164](#), [182](#), [218](#), [220](#), [222](#), [242](#), [247](#), [250](#), [254](#), [255](#), [412](#), [436](#)
- DWARF Version 6, [iii](#), [22](#), [57](#), [161](#), [436](#)
- dwo\_id, [70](#), [211](#), [415](#), [420](#)
- Dylan, [64](#)
- dynamic number of array
  - dimensions, [21](#)
  
- elemental attribute, [80](#), [221](#)
- elemental property of a subroutine, [20](#)
- elements of breg subrange type, [19](#)
- empty location description, [39](#)
- encoding attribute, [106](#), [107](#), [220](#), [238](#)
- encoding of base type, [20](#)
- end-of-list
  - of location list, [44](#)
  - of range list, [54](#)
- end-of-list entry
  - in location list, [237](#)
- end\_sequence, [157](#), [159](#), [171](#)
- endianness attribute, [102](#), [106](#), [221](#)
- endianness of data, [20](#)
- entity, [15](#)
- entry address, [56](#)
- entry address of a scope, [20](#)
- entry PC address, [56](#)
  
- entry PC attribute, [220](#)
  - and abstract instance, [85](#)
  - for catch block, [97](#)
  - for inlined subprogram, [86](#)
  - for lexical block, [95](#)
  - for module initialization, [74](#)
  - for subroutine, [81](#)
  - for try block, [97](#)
  - for with statement, [96](#)
- entry pc attribute, [68](#)
  - and abstract instance, [85](#)
- entry point entry, [78](#)
- enum class, *see* type-safe enumeration
- enum, Rust, [329](#)
- enumeration class attribute, [221](#)
- enumeration literal, *see* enumeration entry
  - enumeration literal value, [19](#)
- enumeration stride (dimension of array type), [18](#)
- enumeration type entry, [130](#)
  - as array dimension, [116](#), [130](#)
- enumerator entry, [130](#)
- epilogue, [179](#), [188](#)
- epilogue begin, [171](#)
- epilogue code, [179](#)
- epilogue end, [171](#)
- epilogue\_begin, [158](#), [159](#), [167](#), [169](#), [171](#)
- error value, [192](#)
- exception thrown, *see* thrown type entry
- explicit attribute, [124](#), [221](#)
- explicit property of member function, [20](#)
- export symbols (of structure, class or union) attribute, [118](#)
- export symbols attribute, [74](#), [222](#)
- export symbols of a namespace, [20](#)



## Index

- export symbols of a structure, union or class, [20](#)
- exprloc, *see also* exprloc class
- exprval class, [23](#), [57](#), [186–188](#), [218–222](#), [225](#), [231](#)
- extended opcodes in line number program, [171](#)
- extended type (Java), *see* inheritance entry
- extensibility, *see* vendor extensibility
- extension attribute, [74](#), [221](#)
- external attribute, [78](#), [99](#), [220](#)
- external subroutine, [20](#)
- external variable, [20](#)
  
- file, [157](#)
- file containing call site of non-inlined call, [18](#)
- file containing declaration, [50](#)
- file containing inlined subroutine call, [18](#)
- file containing source declaration, [19](#)
- file type entry, [136](#)
- file\_name\_format\_count, [163](#), [339](#)
- file\_name\_format\_table, [163](#), [339](#)
- file\_names, [163–165](#), [339](#)
- file\_names\_count, [163](#), [339](#)
- flag, *see also* flag class, [48](#), [49](#), [68](#), [78](#), [80](#), [86](#), [99–101](#), [123](#), [130](#), [131](#), [134](#)
- flag class, [24](#), [74](#), [78](#), [80](#), [91](#), [93](#), [102](#), [195](#), [219–223](#), [225](#), [230](#), [231](#), [245](#)
- foreign\_type\_unit\_count, [150](#), [151](#)
- formal parameter, [82](#)
- formal parameter entry, [99](#), [129](#)
  - in catch block, [97](#)
  - with default value, [101](#)
- formal type parameter, *see* template type parameter entry
- FORTRAN, [9](#)
  
- Fortran, [1](#), [32](#), [64](#), [68](#), [73](#), [76](#), [77](#), [79](#), [80](#), [82](#), [103](#), [132](#), [306](#), [310–312](#), [376](#), [391](#), [394](#)
  - common block, [82](#), [102](#)
  - main program, [79](#)
  - module (Fortran 90), [73](#)
  - use statement, [76](#), [77](#)
- Fortran 2003, [109](#), [132](#), [369](#)
- Fortran 90, [12](#), [73](#), [136](#), [138](#), [304](#), [306](#), [316](#)
- Fortran 90 array, [138](#)
- Fortran array, [304](#)
- Fortran array example, [306](#)
- Fortran example, [394](#)
- frame base attribute, [82](#), [220](#)
- frame description entry, [181](#)
- friend attribute, [122](#), [220](#)
- friend entry, [122](#)
- friend relationship, [20](#)
- full compilation unit, [62](#)
- function entry, *see* subroutine entry
- function template instantiation, [84](#)
- fundamental type, *see* base type entry
  
- generic type, [27](#), [28](#), [30–33](#), [35](#), [37](#)
- global namespace, [74](#), [75](#), *see* namespace (C++), global
- GNU C, [112](#)
- Go, [64](#)
  
- Haskell, [64](#)
- header\_length, [160](#), [206](#), [339](#)
- hidden indirection, *see* data location attribute
- high PC attribute, [52](#), [53](#), [63](#), [74](#), [81](#), [86](#), [95–97](#), [218](#), [351](#)
  - and abstract instance, [85](#)
- high user attribute encoding, [223](#)
- High-Level Shading Language, [64](#)
  
- identifier case attribute, [67](#), [220](#)
- identifier case rule, [20](#)

## Index

- identifier names, [51](#), [58](#), [67](#)
- IEEE 754R decimal floating-point number, [108](#)
- immutable type, [113](#)
- implementing type (Java), *see* inheritance entry
- implicit location description, [40](#)
- implicit pointer example, [366](#), [368](#)
- import attribute, [75](#), [76](#), [78](#), [219](#)
- imported declaration, [20](#)
- imported declaration entry, [75](#)
- imported module attribute, [76](#)
- imported module entry, [76](#)
- imported unit, [20](#)
- imported unit entry, [62](#), [78](#)
- include\_index, [165](#)
- incomplete declaration, [49](#)
- incomplete structure/union/class, [119](#)
- incomplete type, [115](#), [119](#)
- incomplete type (Ada), [112](#)
- incomplete, non-defining, or separate declaration corresponding to a declaration, [22](#)
- incomplete, non-defining, or separate entity declaration, [19](#)
- index attribute, [147](#)
- indirection to actual data, [19](#)
- inheritance entry, [121](#)
- inherited member location, [19](#)
- initial length, [141](#), [149](#), [154](#), [160](#), [182](#), [183](#), [192](#), [193](#), [206](#), [209–211](#), [246](#), [253–255](#)
- initial length field, *see* initial length
- initial\_instructions, [183](#), [189](#)
- initial\_length, [209](#)
- initial\_location, [182](#), [183](#), [189](#), [195](#)
- inline attribute, [84](#), [85](#), [219](#), [244](#)
- inline instances of inline subprograms, [17](#)
- inline namespace, *see also* export symbols attribute, [75](#)
- inlined call location attributes, [86](#)
- inlined subprogram call examples, [347](#)
- inlined subprogram entry, [78](#), [86](#) in concrete instance, [86](#)
- inlined subroutine, [20](#)
- instructions, [183](#)
- integer constant, [50](#), [58](#), [79](#), [84](#), [86](#), [106](#), [107](#), [109](#), [111](#), [115](#), [123](#), [124](#)
- interface type entry, [120](#)
- Internet, [166](#)
- is optional attribute, [101](#), [219](#)
- is\_stmt, [157](#), [159](#), [161](#), [169](#)
- is\_supplementary, [203](#)
- isa, [158](#), [159](#), [171](#)
- ISO 10646 (Fortran string kind), [109](#)
- ISO 10646 character set standard, [109](#), [132](#), [229](#), [369](#)
- ISO-defined language names, [63](#), [241](#)
- ISO/IEC 10646-1:1993 character, [108](#)
- ISO/IEC 10646-1:1993 character (UCS-4), [108](#)
- ISO/IEC 646:1991 character, [108](#)
- ISO\_10646 (Fortran string kind), [132](#)
- Java, [12](#), [64](#), [118](#), [120](#), [121](#)
- Julia, [64](#)
- Kotlin, [64](#)
- label entry, [96](#)
- language attribute, [72](#)
- language attribute (Version 5), [218](#)
- language attribute, encoding, [241](#)
- language name attribute, [63](#), [116](#), [223](#)
- language name encoding, [241](#)
- language version attribute, [223](#)
- language version encoding schemes, [66](#)



## Index

- LEB128, [159](#), [162](#), [170](#), [224](#), [229](#), [232](#)
  - examples, [233](#)
  - signed, [29](#), [41](#), [169](#), [183](#), [185–187](#), [217](#), [232](#), [233](#)
  - signed, decoding of, [296](#)
  - signed, encoding as, [232](#), [295](#)
  - unsigned, [27–29](#), [34](#), [37](#), [40](#), [42](#), [159](#), [169](#), [171](#), [172](#), [182](#), [183](#), [185–188](#), [212](#), [213](#), [217](#), [224–227](#), [232](#), [233](#), [249](#), [258](#), [260](#)
  - unsigned, decoding of, [295](#)
  - unsigned, encoding as, [232](#), [294](#)
- LEB128 encoding
  - algorithms, [294](#)
  - examples, [233](#)
- length, [182–184](#)
- level-88 condition, COBOL, [129](#)
- lexical block, [38](#), [95](#), [97](#)
- lexical block entry, [95](#)
- lexical blocks, [56](#)
- line, [157](#)
- line containing call site of
  - non-inlined call, [18](#)
- line number information, *see also* statement list attribute
- line number information for unit, [22](#)
- line number of inlined subroutine call, [18](#)
- line number of source declaration, [19](#)
- line number opcodes
  - extended opcode encoding, [248](#)
  - file entry format encoding, [248](#)
  - standard opcode encoding, [247](#)
- line number program, [167](#)
  - extended opcodes, [171](#)
  - special opcodes, [167](#)
  - standard opcodes, [169](#)
- line\_base, [161](#), [168](#), [169](#), [339](#), [340](#)
- line\_range, [161](#), [168](#), [169](#), [339](#), [340](#)
- lineptr, *see also* lineptr class, [287](#)
- lineptr class, [24](#), [72](#), [218](#), [223](#), [225](#), [230](#), [231](#), [292](#)
- linkage name attribute, [58](#), [221](#)
- list of discriminant values, [20](#)
- Little-Endian Base 128, *see* LEB128
- little-endian encoding, *see* endian attribute
- local\_type\_unit\_count, [150](#), [151](#)
- location, [123](#)
- location attribute, [36](#), [51](#), [96](#), [100](#), [102](#), [218](#), [351](#)
  - and abstract instance, [85](#)
- location description, [38](#), *see also* DWARF expression, [121](#), [123](#), [125](#), [132](#), [135](#), [137](#)
  - composite, [39](#)
  - empty, [39](#)
  - implicit, [40](#)
  - memory, [39](#)
  - simple, [39](#)
  - single, [38](#)
  - use in location list, [39](#)
- location list, [39](#), [63](#), [82](#), [83](#), [237](#), [261](#), [288](#), [292](#)
- location list attribute, [220](#)
- location list base attribute, [222](#)
- location lists base, [21](#)
- location of uplevel frame, [22](#)
- location table base attribute, [69](#)
- locdesc class, [24](#), [39](#), [57](#), [218–220](#), [222](#), [226](#), [231](#)
- loclist, *see also* loclist class
- loclist class, [24](#), [39](#), [43](#), [218–220](#), [223](#), [226](#), [231](#), [288](#), [292](#)
- loclistsptr, *see also* loclistsptr class
- loclistsptr class, [24](#), [69](#), [222](#), [223](#), [226](#), [230](#), [231](#)
- lookup
  - by address, [154](#)
  - by name, [142](#)

## Index

- low PC attribute, [52](#), [53](#), [63](#), [74](#), [81](#), [86](#),  
[95–97](#), [218](#), [351](#)
  - and abstract instance, [85](#)
- low user attribute encoding, [223](#)
- lower bound attribute, [134](#), [219](#)
  - default, [134](#), [241](#)
- lower bound of subrange, [21](#)
- macptr, *see also* macptr class, [287](#)
- macptr class, [24](#), [71](#), [222](#), [223](#), [226](#),  
[230](#), [231](#), [292](#)
- macro formal parameter list, [176](#)
- macro information, [172](#)
- macro information attribute, [65](#), [222](#)
- macro information entry types
  - encoding, [249](#)
- macro preprocessor information, [21](#)
- main or starting subprogram, [21](#)
- main subprogram attribute, [68](#), [78](#),  
[221](#)
- mangled names, [51](#), [58](#), [91](#)
- maximum\_operations\_per\_instruction,  
[160](#), [161](#), [168](#), [339](#)
- MD5, [166](#), [257](#), [261](#), [262](#), [399](#), [406](#)
- member entry (data), [122](#), [123](#)
  - as discriminant, [127](#)
- member function entry, [124](#)
- member function example, [335](#)
- member location for pointer to
  - member type, [22](#)
- memory location description, [39](#)
- minimum\_instruction\_length, [160](#),  
[161](#), [168](#), [170](#), [339](#)
- MIPS instruction set architecture, [155](#)
- Modula-2, [47](#), [64](#), [73](#), [96](#)
  - definition module, [73](#)
- Modula-3, [64](#)
- module entry, [74](#)
- module priority, [21](#)
- Mojo, [64](#)
- mutable attribute, [123](#), [221](#)
- mutable property of member data, [21](#)
- name attribute, [51](#), [59](#), [63](#), [67](#), [74](#), [76](#),  
[88](#), [95](#), [96](#), [99](#), [102](#), [103](#), [106](#),  
[112](#), [113](#), [115](#), [118](#), [120](#), [122](#),  
[129–133](#), [135–137](#), [218](#), [258](#),  
[260](#)
- name index, [142](#)
  - case folding, [151](#)
- name list item attribute, [220](#)
- name of declaration, [21](#)
- name\_count, [150–152](#)
- namelist entry, [103](#)
- namelist item, [21](#)
- namelist item attribute, [103](#)
- namelist item entry, [103](#)
- names
  - identifier, [51](#)
  - mangled, [58](#)
- namespace (C++), [74](#)
  - alias, [76](#)
  - example, [331](#)
  - global, [75](#)
  - unnamed, [75](#)
  - using declaration, [75–77](#)
  - using directive, [77](#)
- namespace alias, [20](#)
- namespace declaration entry, [74](#)
- namespace extension entry, [74](#)
- namespace using declaration, [20](#)
- namespace using directive, [20](#)
- nested abstract instance, [85](#)
- nested concrete inline instance, [87](#)
- non-constant parameter flag, [22](#)
- non-contiguous address ranges, [53](#)
- non-contiguous range of code
  - addresses, [21](#)
- non-default alignment, [17](#)
- non-defining declaration, [49](#)
- noreturn attribute, [21](#), [80](#), [222](#)

## Index

- numerator of rational scale factor, [22](#), [222](#)
- object (*this*, *self*) pointer of member function, [21](#)
- object file linkage name of an entity, [21](#)
- object pointer attribute, [125](#), [221](#)
- Objective C, [64](#), [125](#)
- Objective C++, [64](#)
- objects or types that are not actually declared in the source, [17](#)
- OCaml, [64](#)
- `offset_entry_count`, [254–256](#)
- `offset_size_flag`, [173](#), [175](#), [178](#), [200–202](#), [380](#), [381](#)
- `op_index`, [157–161](#), [167–170](#), [172](#)
- `opcode_base`, [161](#), [162](#), [168](#), [339](#)
- `opcode_operands_table`, [173](#), [174](#)
- `opcode_operands_table_flag`, [173](#), [380](#), [381](#)
- OpenCL C, [9](#), [65](#), [166](#)
- OpenCL C++, [65](#)
- OpenGL ES Shading Language, [64](#)
- OpenGL Shading Language, [64](#)
- operation advance, [168](#), [169](#), [340](#)
- operation pointer, [158](#), [161](#), [167](#), [168](#)
- optional parameter, [20](#), [101](#)
- ordering attribute, [102](#), [218](#)
- out-of-line instance, [90](#), *see also* concrete out-of-line instance
- out-of-line instances of inline subprograms, [17](#)
- package files, [198](#)
- packed qualified type entry, [113](#)
- packed type entry, [113](#)
- padding, [149](#), [183](#), [184](#), [200](#), [253](#)
- parameter, *see this* parameter, *see* formal parameter entry, *see* macro formal parameter list, *see* optional parameter attribute, *see* template type parameter entry, *see* template value parameter entry, *see* unspecified parameters entry, *see* variable parameter attribute
- parameter entry, [18](#)
- partial compilation unit, [62](#), [63](#)
- Pascal, [65](#), [96](#), [113](#), [118](#), [133](#), [136](#), [326](#), [327](#), [329](#)
- Pascal example, [321](#)
- path name of compilation source, [21](#)
- picture string attribute, [221](#)
- picture string for numeric string type, [21](#)
- PL/I, [65](#)
- pointer or reference types, [17](#)
- pointer qualified type entry, [113](#)
- pointer to member, [135](#)
- pointer to member entry, [135](#)
- pointer to member type, [135](#), [136](#)
- pointer to member type entry, [135](#)
- pointer type entry, [114](#)
- previous namespace extension or original namespace, [20](#)
- primitive data types of compilation unit, [17](#)
- priority attribute, [74](#), [220](#)
- producer attribute, [67](#), [219](#)
- PROGRAM statement, [68](#)
- programming language name, [21](#)
- programming language version, [21](#)
- prologue, [179](#)
- prologue end, [170](#), [171](#)
- `prologue_end`, [158](#), [159](#), [167](#), [169](#), [170](#)
- prototyped attribute, [80](#), [131](#), [219](#)
- pure attribute, [80](#), [221](#)
- pure property of a subroutine, [21](#)
- Python, [65](#)

## Index

- range list, [63](#), [98](#), [252](#), [253](#), [288](#), [292](#)
- range list base
  - encoding, [222](#)
- ranges attribute, [52](#), [53](#), [63](#), [74](#), [81](#), [86](#), [95–97](#), [221](#)
  - and abstract instance, [85](#)
- ranges lists, [22](#)
- ranges table base attribute, [69](#)
- rank attribute, [221](#)
- recursive attribute, [80](#), [221](#)
- recursive property of a subroutine, [21](#)
- reduced scope of declaration, [22](#)
- reference, *see also* reference class, [50](#), [57](#), [68](#), [75](#), [76](#), [78](#), [82](#), [83](#), [100](#), [103](#), [113](#), [115](#), [119](#), [125](#), [127](#), [132](#), [135](#)
- reference attribute, [222](#)
- reference class, [24](#), [90](#), [93](#), [101](#), [102](#), [132](#), [218–222](#), [227](#), [231](#), [245](#)
- reference qualified type entry, [113](#)
- reference type, [114](#)
- reference type entry, [113](#)
- reference type entry, lvalue, *see* reference type entry
- reference type entry, rvalue, *see* rvalue reference type entry
- renamed declaration, *see* imported declaration entry
- RenderScript, [65](#)
- reserved target address, [26](#)
- reserved values
  - error, [192](#)
  - initial length, [192](#)
- restrict qualified type, [113](#)
- restricted type entry, [113](#)
- return address attribute, [82](#), [219](#)
  - and abstract instance, [85](#)
- return address from a call, [18](#)
- return type of subroutine, [81](#)
- return\_address\_register, [183](#)
- rnglist, *see also* rnglist class
- rnglist class, [24](#), [53](#), [219](#), [221](#), [223](#), [227](#), [231](#), [288](#), [292](#)
- rnglistsptr, *see also* rnglistsptr class
- rnglistsptr class, [24](#), [69](#), [98](#), [222](#), [223](#), [227](#), [230](#), [231](#)
- Ruby, [65](#)
- Rust, [65](#), [112](#), [329](#), [330](#)
- rvalue reference qualified type entry, [113](#)
- rvalue reference type entry, [113](#)
- rvaluereference attribute, [222](#)
- sbyte, [141](#), [161](#), [257](#)
- scalar coarray, *see* coarray
- scale factor for fixed-point type, [22](#)
- scaled encodings, [109](#)
  - binary, [109](#)
  - composition of, [110](#)
  - decimal, [109](#)
  - floating-point, [110](#)
  - rational, [110](#)
- section group, [387–394](#), [396](#), [398](#), [409](#), [410](#)
  - name, [389](#)
- section length, [141](#)
  - use in headers, [206](#)
- section offset, [11](#), [141](#), [210–212](#), [246](#)
  - alignment of, [256](#), [257](#)
  - in .debug\_aranges header, [154](#), [246](#)
  - in .debug\_info header, [210–212](#)
  - in class lineptr value, [225](#)
  - in class loclist value, [226](#)
  - in class loclistsptr, [226](#)
  - in class macptr value, [226](#)
  - in class reference value, [227](#)
  - in class rnglist value, [227](#)
  - in class rnglistsptr, [227](#)
  - in class string value, [229](#)
  - in FDE header, [183](#)
  - in macro information attribute, [65](#)

## Index

- in statement list attribute, 65
- use in headers, 206
- section\_count, 200
- self pointer attribute, *see* object pointer attribute
- set type entry, 133
- shared qualified type entry, 113
- sibling attribute, 25, 50, 218
- signature attribute, 221
- signed LEB128, *see* LEB128, signed
- simple location description, 39
- single location description, 38
- size of an address, 26, *see also*
  - address\_size, 27, 30, 31, 132, 154, 160, 211, 246, 247
- skeleton compilation unit, 69
- SLEB128, 232
- slot\_count, 200
- small attribute, 110, 221
- special opcodes in line number program, 167
- specialized .debug\_line.dwo section, 424
- specialized line number table, 72, 196, 413
- specification attribute, 75, 85, 119, 125, 126, 220
- split DWARF object file, 69, 70, 72, 144, 148, 195, 198, 223, 285, 412, 417, 420, 421, 423
  - example, 417
  - object file name, 20
- split DWARF object file name encoding, 222
- split DWARF object file name attribute, 69
- split type unit, 71
- standard opcodes in line number program, 169
- standard\_opcode\_lengths, 162, 339
- start scope attribute, 97, 219
  - and abstract instance, 85
- statement list attribute, 65, 72, 218
- static link attribute, 220
- stride attribute, *see* bit stride attribute or byte stride attribute
- string, *see also* string class
- string class, 24, 51, 91, 101, 218, 219, 221, 222, 228, 230–232
- string length attribute, 132, 219
  - size of length, 221
  - size of length data, 132
- string length of string type, 22
  - size of, 22
- string length size attribute, 132
- string offset section attribute, 68
- string offsets attribute, 222
- string offsets base attribute, 72
- string offsets information, 287
- string offsets information for unit, 22
- string type entry, 132
- stroffsetsptr, *see also* stroffsetsptr class
- stroffsetsptr class, 25, 68, 72, 222, 223, 230, 231
- structure type entry, 118
- subprogram called, 18
- subprogram entry, 78
  - as member function, 124
  - use for template instantiation, 84
  - use in inlined subprogram, 84
- subrange stride (dimension of array type), 18
- subrange type entry, 133
  - as array dimension, 116
- subroutine call site summary
  - attributes, 80
- subroutine entry, 78
- subroutine formal parameters, 82
- subroutine frame base address, 20
- subroutine or subroutine type, 17
- subroutine prototype, 21

## Index

- subroutine return address save
  - location, 21
- subroutine type entry, 131
- sup\_checksum, 203
- sup\_checksum\_len, 203
- sup\_filename, 203
- supplementary object file, 8, 17, 24, 175, 178, 203, 204, 228, 229, 284, 285
- Swift, 65
- SYCL, 65
- tag, 15
- tag names, *see* debugging information entry
  - list of, 23
- target address, 183
  - reserved, 26
- target subroutine of trampoline, 22
- template alias entry, 137
- template alias example, 362, 363
- template alias example 1, 363
- template alias example 2, 364
- template instantiation, 58
  - and special compilation unit, 127
  - function, 84
- template type parameter entry, 59
- template value parameter, 19
- template value parameter entry, 59
- Tensor (array) type, 22
- tensor attribute, 223
- this parameter, 48, 91
- this pointer attribute, *see* object pointer attribute
- thread scaled attribute, 221
- thread-local storage, 32
- threads scaled attribute, 134
- thrown exception, *see* thrown type entry
- thrown type entry, 83
- trampoline (subprogram) entry, 90
- trampoline attribute, 90, 221
- try block, 97
- try block entry, 97
- try/catch blocks, 56
- type
  - of call site, 22
  - of declaration, 22
  - of string type components, 22
  - of subroutine return, 22
- type attribute, 47, 59, 81, 83, 96, 113, 115, 116, 121, 122, 127, 129–131, 133–136, 220
  - of call site entry, 93
  - of string type entry, 132
- type modifier, *see* atomic type entry, *see* constant type entry, *see* packed type entry, *see* pointer type entry, *see* reference type entry, *see* restricted type entry, *see* shared type entry, *see* volatile type entry
- type modifier entry, 113
- type safe enumeration definition, 20
- type safe enumeration types, 130
- type signature, 22, 24, 119, 212, 228, 257, 398, 399
  - computation, 257
  - computation grammar, 406
  - example computation, 398
- type unit, 71, *see also* compilation unit, 73, 119, 208, 211, 212, 228, 257, 260, 261, 399, 410
  - specialized `.debug_line.dwo` section in, 424
- type unit entry, 72
- type unit set, 199
- type unit signature, 70, 430
- type-safe enumeration, 359
- type\_offset, 212
- type\_signature, 212
- typedef entry, 115



## Index

- ubyte, [141](#), [150](#), [154](#), [158–162](#), [167](#),  
[173](#), [182–185](#), [203](#), [210](#), [211](#),  
[246](#), [254](#), [255](#), [257](#)
- UCS character, [108](#)
- uhalf, [141](#), [149](#), [154](#), [160](#), [170](#), [173](#), [185](#),  
[200](#), [203](#), [209–211](#), [246](#),  
[253–255](#), [257](#)
- ULEB128, [174](#), [175](#), [232](#)
- unallocated variable, [100](#)
- Unicode, [109](#), [151](#), [229](#), [358](#), *see also*  
UTF-8
- Unified Parallel C, *see* UPC
- union type entry, [118](#)
- unit, *see* compilation unit
- unit containing main or starting  
subprogram, [21](#)
- unit header unit type encodings, [209](#)
- unit\_count, [200](#)
- unit\_length, [149](#), [154](#), [160](#), [209–211](#),  
[246](#), [253–255](#), [339](#)
- unit\_type, [9](#), [209–211](#)
- unnamed namespace, *see* namespace  
(C++), unnamed
- unsigned LEB128, *see* LEB128,  
unsigned
- unspecified parameters entry, [82](#), [131](#)  
in catch block, [97](#)
- unspecified type entry, [112](#)
- unwind, *see* virtual unwind
- UPC, [22](#), [65](#), [113](#), [114](#), [134](#)
- uplevel address, *see* static link  
attribute
- upper bound attribute, [134](#), [219](#)  
default unknown, [134](#)
- upper bound of subrange, [22](#)
- use location attribute, [135](#)
- use statement, *see* Fortran, use  
statement
- use UTF8 attribute, [22](#), [68](#), [220](#), [229](#)
- using declaration, *see* namespace  
(C++), using declaration
- using directive, *see* namespace (C++),  
using directive
- UTF character, [108](#)
- UTF-8, [12](#), [22](#), [68](#), [182](#), [220](#), [229](#)
- uword, [141](#), [150](#), [185](#), [200](#), [254](#), [255](#),  
[257](#)
- value pointed to by an argument, [18](#)
- variable entry, [99](#)
- variable length data, *see* LEB128
- variable parameter attribute, [100](#), [220](#)
- variant entry, [128](#)
- variant part entry, [127](#)
- vendor extensibility, [191](#)
- vendor id, [192](#)
- vendor specific extensions, *see*  
vendor extensibility
- version, [149](#), [154](#), [160](#), [182](#), [200](#), [203](#),  
[209–211](#), [246](#), [253–255](#), [339](#)
- version number
  - address lookup table, [154](#)
  - address range table, [246](#)
  - address table, [254](#)
  - call frame information, [182](#), [251](#)
  - compilation unit, [209](#), [210](#)
  - CU index information, [200](#)
  - line number information, [160](#), [247](#)
  - location list table, [255](#)
  - macro information, [249](#)
  - name index table, [149](#), [245](#)
  - range list table, [254](#)
  - string offsets table, [253](#)
  - summary by section, [436](#)
  - TU index information, [200](#)
  - type unit, [211](#)
- virtual function vtable slot, [23](#)
- virtual unwind, [178](#)
- virtuality attribute, [48](#), [220](#)
- virtuality of member function or base  
class, [23](#)
- visibility attribute, [47](#), [219](#)

## Index

visibility of declaration, [23](#)  
void type, *see* unspecified type entry  
volatile qualified type entry, [113](#)  
vtable element location attribute, [125](#),  
[220](#)  
VVM language version encoding  
scheme, [64–66](#)  
VVMPP language version  
encoding scheme, [64, 66](#)  
  
with statement entry, [96](#)  
  
YYYY language version encoding  
scheme, [64–66](#)  
YYYYMM language version  
encoding scheme, [64, 66](#)  
YYYYRR language version encoding  
scheme, [65, 66](#)  
  
Zig, [65](#)