

GNU Offloading and Multi Processing Runtime Library

The GNU OpenMP and OpenACC Implementation

Published by the Free Software Foundation
51 Franklin Street, Fifth Floor
Boston, MA 02110-1301, USA

Copyright © 2006-2025 Free Software Foundation, Inc.

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 the Invariant Sections being “Funding Free Software”, the Front-Cover texts being (a) (see below), and with the Back-Cover Texts being (b) (see below). A copy of the license is included in the section entitled “GNU Free Documentation License”.

(a) The FSF’s Front-Cover Text is:

A GNU Manual

(b) The FSF’s Back-Cover Text is:

You have freedom to copy and modify this GNU Manual, like GNU software. Copies published by the Free Software Foundation raise funds for GNU development.

Short Contents

1	Enabling OpenMP.....	1
2	OpenMP Implementation Status	3
3	OpenMP Runtime Library Routines	15
4	OpenMP Environment Variables	59
5	Enabling OpenACC.....	71
6	OpenACC Runtime Library Routines	73
7	OpenACC Environment Variables	93
8	CUDA Streams Usage.....	95
9	OpenACC Library Interoperability	97
10	OpenACC Profiling Interface	101
11	OpenMP-Implementation Specifics.....	107
12	Offload-Target Specifics.....	113
13	The libgomp ABI.....	119
14	Reporting Bugs	125
	GNU General Public License	127
	GNU Free Documentation License.....	139
	Funding Free Software	147
	Library Index	149

Table of Contents

1	Enabling OpenMP	1
2	OpenMP Implementation Status.....	3
2.1	OpenMP 4.5.....	3
2.2	OpenMP 5.0.....	3
	New features listed in Appendix B of the OpenMP specification ...	3
	Other new OpenMP 5.0 features	5
2.3	OpenMP 5.1.....	5
	New features listed in Appendix B of the OpenMP specification ...	5
	Other new OpenMP 5.1 features	6
2.4	OpenMP 5.2.....	7
	New features listed in Appendix B of the OpenMP specification ...	7
	Other new OpenMP 5.2 features	8
2.5	OpenMP 6.0.....	9
	New features listed in Appendix B of the OpenMP specification ...	9
	Deprecated features, unless listed above	12
	Other new OpenMP 6.0 features	12
2.6	OpenMP Technical Report 14	13
	New features listed in Appendix B of the OpenMP specification ..	13
	Deprecated features, unless listed above	13
3	OpenMP Runtime Library Routines	15
3.1	Thread Team Routines	15
3.1.1	omp_set_num_threads – Set upper team size limit	15
3.1.2	omp_get_num_threads – Size of the active team.....	15
3.1.3	omp_get_max_threads – Maximum number of threads of parallel region	16
3.1.4	omp_get_thread_num – Current thread ID	16
3.1.5	omp_in_parallel – Whether a parallel region is active....	16
3.1.6	omp_set_dynamic – Enable/disable dynamic teams.....	17
3.1.7	omp_get_dynamic – Dynamic teams setting	17
3.1.8	omp_get_cancellation – Whether cancellation support is enabled	17
3.1.9	omp_set_nested – Enable/disable nested parallel regions..	18
3.1.10	omp_get_nested – Nested parallel regions	18
3.1.11	omp_set_schedule – Set the runtime scheduling method ..	19
3.1.12	omp_get_schedule – Obtain the runtime scheduling method..	19
3.1.13	omp_get_teams_thread_limit – Maximum number of threads imposed by teams	20
3.1.14	omp_get_supported_active_levels – Maximum number of active regions supported	20
3.1.15	omp_set_max_active_levels – Limits the number of active parallel regions.....	20

<code>self_maps</code> clause to <code>requires</code> directive	Y
<code>no_openmp_constructs</code> assumptions clause	N
Restriction for <code>ordered</code> regarding loop-transforming directives	N
<code>apply</code> clause to loop-transforming constructs	N
Non-constant values in the <code>sizes</code> clause	N
<code>fuse</code> loop-transformation construct	N
<code>interchange</code> loop-transformation construct	N
<code>reverse</code> loop-transformation construct	N
<code>split</code> loop-transformation construct	N
<code>stripe</code> loop-transformation construct	N
<code>tile</code> permitting association of grid and inter-tile loops	N
<code>strict</code> modifier keyword to <code>num_threads</code>	N
<code>safesync</code> clause to the <code>parallel</code> construct	N
<code>omp_curr_progress_width</code> identifier	N
<code>omp_get_max_progress_width</code> runtime routine	N
Lifted restrictions on <code>order(concurrent)</code> and, hence, the <code>loop</code> construct	N
<code>atomic</code> permitted in a construct with <code>order(concurrent)</code>	N
Lifted restrictions on not-strictly-nested regions with <code>order(concurrent)</code>	N
<code>workdistribute</code> directive for Fortran	N
Fortran DO CONCURRENT as associated loop in a <code>loop</code> construct	N
New <code>task_iteration</code> directive inside <code>taskloop</code>	N
<code>threadset</code> clause in task-generating constructs	N
New <code>priority</code> clause to <code>target</code> , <code>target_enter_data</code> , <code>target_data</code> , <code>target_exit_data</code> and <code>target_update</code>	N
New <code>device_type</code> clause to the <code>target</code> directive	N
<code>target_data</code> as composite construct	N
<code>nowait</code> clause with reverse-offload <code>target</code> directives	N
Extended <i>prefer-type</i> modifier to <code>init</code> clause	Y
Boolean argument to <code>nowait</code> and <code>nogroup</code> may be non constant	N
<code>memscope</code> clause to <code>atomic</code> and <code>flush</code>	N
New <code>transparent</code> clause for multi-generational task-dependence graphs	N
The <code>cancel</code> construct now completes tasks with unfulfilled events	N
<code>omp_fulfill_event</code> routine was restricted regarding fulfillment of event variables	N
Added rule for compound-directive names, permitting many more combinations	N
<code>omp_is_free_agent</code> and <code>omp_ancestor_is_free_agent</code> routines	N

3.7.6 `omp_target_memcpy_async` – Copy data between devices asynchronously

Description:

This routine copies asynchronously *length* of bytes of data from the device identified by device number *src_device_num* to device *dst_device_num*. The data is copied from the source device from the address provided by *src*, shifted by the offset of *src_offset* bytes, to the destination device's *dst* address shifted by *dst_offset*. Task dependence is expressed by passing an array of depend objects to *depobj_list*, where the number of array elements is passed as *depobj_count*; if the count is zero, the *depobj_list* argument is ignored. In C++ and Fortran, the *depobj_list* argument can also be omitted in that case. The routine returns zero if the copying process has successfully been started and non-zero otherwise.

Running this routine in a `target` region except on the initial device is not supported.

C/C++

Prototype:

```
int omp_target_memcpy_async(void *dst,
    const void *src,
    size_t length,
    size_t dst_offset,
    size_t src_offset,
    int dst_device_num,
    int src_device_num,
    int depobj_count,
    omp_depend_t *depobj_list)
```

Fortran:

Interface:

```
integer(c_int) function omp_target_memcpy_async( &
    dst, src, length, dst_offset, src_offset, &
    dst_device_num, src_device_num, &
    depobj_count, depobj_list) bind(C)
use, intrinsic :: iso_c_binding, only: c_ptr,
c_size_t, c_int
type(c_ptr), value :: dst, src
integer(c_size_t), value :: length, dst_offset,
src_offset
integer(c_int), value :: dst_device_num, src_
device_num, depobj_count
integer(omp_depend_kind), optional :: depobj_
list(*)
```

See also: Section 3.7.5 [`omp_target_memcpy`], page 33, Section 3.7.8 [`omp_target_memcpy_rect_async`], page 36,

Reference: OpenMP specification v5.1 (<https://www.openmp.org>), Section 3.8.7

- `$<priority>` is an optional priority for the worker threads of a thread pool according to `pthread_setschedparam`. In case a priority value is omitted, then a worker thread inherits the priority of the OpenMP primary thread that created it. The priority of the worker thread is not changed after creation, even if a new OpenMP primary thread using the worker has a different priority.
- `@<scheduler-name>` is the scheduler instance name according to the RTEMS application configuration.

In case no thread pool configuration is specified for a scheduler instance, then each OpenMP primary thread of this scheduler instance uses its own dynamically allocated thread pool. To limit the worker thread count of the thread pools, each OpenMP primary thread must call `omp_set_num_threads`.

Example: Lets suppose we have three scheduler instances `I0`, `WRK0`, and `WRK1` with `GOMP_RTEMS_THREAD_POOLS` set to `"1@WRK0:3$4@WRK1"`. Then there are no thread pool restrictions for scheduler instance `I0`. In the scheduler instance `WRK0` there is one thread pool available. Since no priority is specified for this scheduler instance, the worker thread inherits the priority of the OpenMP primary thread that created it. In the scheduler instance `WRK1` there are three thread pools available and their worker threads run at priority four.

5 Enabling OpenACC

To activate the OpenACC extensions for C/C++ and Fortran, the compile-time flag `-fopenacc` must be specified. This enables the OpenACC directive `#pragma acc` in C/C++ and, in Fortran, the `!$acc` sentinel in free source form and the `c$acc`, `*$acc` and `!$acc` sentinels in fixed source form. The flag also arranges for automatic linking of the OpenACC runtime library (Chapter 6 [OpenACC Runtime Library Routines], page 73).

See <https://gcc.gnu.org/wiki/OpenACC> for more information.

A complete description of all OpenACC directives accepted may be found in the OpenACC (<https://www.openacc.org>) Application Programming Interface manual, version 2.6.


```

Interface:      type(*), dimension(..) :: a
                 subroutine acc_create(a, len)
                 type(*), dimension(..) :: a
                 integer len
Interface:      subroutine acc_create_async(a, async)
                 type(*), dimension(..) :: a
                 integer(acc_handle_kind) :: async
Interface:      subroutine acc_create_async(a, len, async)
                 type(*), dimension(..) :: a
                 integer len
                 integer(acc_handle_kind) :: async

```

Reference: OpenACC specification v2.6 (<https://www.openacc.org>), section 3.2.21.

6.21 acc_present_or_create – If the data is not present on the device, allocate device memory and map it to host memory.

Description

This function tests if the host data specified by *a* and of length *len* is present or not. If it is not present, device memory is allocated and mapped to host memory. In C/C++, the device address of the newly allocated device memory is returned.

In Fortran, two (2) forms are supported. In the first form, *a* specifies a contiguous array section. The second form *a* specifies a variable or array element and *len* specifies the length in bytes.

Note that `acc_present_or_create` and `acc_pcreate` exist for backward compatibility with OpenACC 2.0; use Section 6.20 [`acc_create`], page 80, instead.

C/C++:

```

Prototype:      void *acc_present_or_create(h_void *a, size_t len)
Prototype:      void *acc_pcreate(h_void *a, size_t len)

```

Fortran:

```

Interface:      subroutine acc_present_or_create(a)
                 type(*), dimension(..) :: a
Interface:      subroutine acc_present_or_create(a, len)
                 type(*), dimension(..) :: a
                 integer len
Interface:      subroutine acc_pcreate(a)
                 type(*), dimension(..) :: a
Interface:      subroutine acc_pcreate(a, len)
                 type(*), dimension(..) :: a
                 integer len

```

Reference: OpenACC specification v2.6 (<https://www.openacc.org>), section 3.2.21.

6.22 `acc_copyout` – Copy device memory to host memory.

Description

This function copies mapped device memory to host memory which is specified by host address *a* for a length *len* bytes in C/C++.

In Fortran, two (2) forms are supported. In the first form, *a* specifies a contiguous array section. The second form *a* specifies a variable or array element and *len* specifies the length in bytes.

C/C++:

```
Prototype:      acc_copyout(h_void *a, size_t len);
Prototype:      acc_copyout_async(h_void *a, size_t len, int async);
Prototype:      acc_copyout_finalize(h_void *a, size_t len);
Prototype:      acc_copyout_finalize_async(h_void *a, size_t len,
                                           int async);
```

Fortran:

```
Interface:      subroutine acc_copyout(a)
                  type(*), dimension(..) :: a
Interface:      subroutine acc_copyout(a, len)
                  type(*), dimension(..) :: a
                  integer len
Interface:      subroutine acc_copyout_async(a, async)
                  type(*), dimension(..) :: a
                  integer(acc_handle_kind) :: async
Interface:      subroutine acc_copyout_async(a, len, async)
                  type(*), dimension(..) :: a
                  integer len
                  integer(acc_handle_kind) :: async
Interface:      subroutine acc_copyout_finalize(a)
                  type(*), dimension(..) :: a
Interface:      subroutine acc_copyout_finalize(a, len)
                  type(*), dimension(..) :: a
                  integer len
Interface:      subroutine acc_copyout_finalize_async(a, async)
                  type(*), dimension(..) :: a
                  integer(acc_handle_kind) :: async
Interface:      subroutine acc_copyout_finalize_async(a, len,
                                                         async)
                  type(*), dimension(..) :: a
                  integer len
                  integer(acc_handle_kind) :: async
```

Reference: OpenACC specification v2.6 (<https://www.openacc.org>), section 3.2.22.

6.23 acc_delete – Free device memory.

Description

This function frees previously allocated device memory specified by the device address *a* and the length of *len* bytes.

In Fortran, two (2) forms are supported. In the first form, *a* specifies a contiguous array section. The second form *a* specifies a variable or array element and *len* specifies the length in bytes.

C/C++:

```

Prototype:      acc_delete(h_void *a, size_t len);
Prototype:      acc_delete_async(h_void *a, size_t len, int async);
Prototype:      acc_delete_finalize(h_void *a, size_t len);
Prototype:      acc_delete_finalize_async(h_void *a, size_t len,
                                           int async);

```

Fortran:

```

Interface:      subroutine acc_delete(a)
                   type(*), dimension(..) :: a
Interface:      subroutine acc_delete(a, len)
                   type(*), dimension(..) :: a
                   integer len
Interface:      subroutine acc_delete_async(a, async)
                   type(*), dimension(..) :: a
                   integer(acc_handle_kind) :: async
Interface:      subroutine acc_delete_async(a, len, async)
                   type(*), dimension(..) :: a
                   integer len
                   integer(acc_handle_kind) :: async
Interface:      subroutine acc_delete_finalize(a)
                   type(*), dimension(..) :: a
Interface:      subroutine acc_delete_finalize(a, len)
                   type(*), dimension(..) :: a
                   integer len
Interface:      subroutine acc_delete_finalize_async(a, async)
                   type(*), dimension(..) :: a
                   integer(acc_handle_kind) :: async
Interface:      subroutine acc_delete_finalize_async(a, len, async)
                   type(*), dimension(..) :: a
                   integer len
                   integer(acc_handle_kind) :: async

```

Reference: OpenACC specification v2.6 (<https://www.openacc.org>), section 3.2.23.

6.40 `acc_prof_register` – Register callbacks.

Description:

This function registers callbacks.

C/C++:

Prototype: `void acc_prof_register (acc_event_t, acc_prof_callback, acc_register_t);`

See also: Chapter 10 [OpenACC Profiling Interface], page 101,

Reference: OpenACC specification v2.6 (<https://www.openacc.org>), section 5.3.

6.41 `acc_prof_unregister` – Unregister callbacks.

Description:

This function unregisters callbacks.

C/C++:

Prototype: `void acc_prof_unregister (acc_event_t, acc_prof_callback, acc_register_t);`

See also: Chapter 10 [OpenACC Profiling Interface], page 101,

Reference: OpenACC specification v2.6 (<https://www.openacc.org>), section 5.3.

6.42 `acc_prof_lookup` – Obtain inquiry functions.

Description:

Function to obtain inquiry functions.

C/C++:

Prototype: `acc_query_fn acc_prof_lookup (const char *);`

See also: Chapter 10 [OpenACC Profiling Interface], page 101,

Reference: OpenACC specification v2.6 (<https://www.openacc.org>), section 5.3.

6.43 `acc_register_library` – Library registration.

Description:

Function for library registration.

C/C++:

Prototype: `void acc_register_library (acc_prof_reg, acc_prof_reg, acc_prof_lookup_func);`

See also: Chapter 10 [OpenACC Profiling Interface], page 101, Section 7.3 [ACC_PROFLIB], page 93,

Reference: OpenACC specification v2.6 (<https://www.openacc.org>), section 5.3.

7 OpenACC Environment Variables

The variables `ACC_DEVICE_TYPE` and `ACC_DEVICE_NUM` are defined by section 4 of the OpenACC specification in version 2.0. The variable `ACC_PROFLIB` is defined by section 4 of the OpenACC specification in version 2.6.

7.1 `ACC_DEVICE_TYPE`

Description:

Control the default device type to use when executing compute regions. If unset, the code can be run on any device type, favoring a non-host device type.

Supported values in GCC (if compiled in) are

- `host`
- `nvidia`
- `radeon`

Reference: OpenACC specification v2.6 (<https://www.openacc.org>), section 4.1.

7.2 `ACC_DEVICE_NUM`

Description:

Control which device, identified by device number, is the default device. The value must be a nonnegative integer less than the number of devices. If unset, device number zero is used.

Reference: OpenACC specification v2.6 (<https://www.openacc.org>), section 4.2.

7.3 `ACC_PROFLIB`

Description:

Semicolon-separated list of dynamic libraries that are loaded as profiling libraries. Each library must provide at least the `acc_register_library` routine. Each library file is found as described by the documentation of `dlopen` of your operating system.

See also: Section 6.43 [`acc_register_library`], page 91, Chapter 10 [OpenACC Profiling Interface], page 101,

Reference: OpenACC specification v2.6 (<https://www.openacc.org>), section 4.3.

8 CUDA Streams Usage

This applies to the `nvptx` plugin only.

The library provides elements that perform asynchronous movement of data and asynchronous operation of computing constructs. This asynchronous functionality is implemented by making use of CUDA streams¹.

The primary means by that the asynchronous functionality is accessed is through the use of those OpenACC directives which make use of the `async` and `wait` clauses. When the `async` clause is first used with a directive, it creates a CUDA stream. If an `async-argument` is used with the `async` clause, then the stream is associated with the specified `async-argument`.

Following the creation of an association between a CUDA stream and the `async-argument` of an `async` clause, both the `wait` clause and the `wait` directive can be used. When either the clause or directive is used after stream creation, it creates a rendezvous point whereby execution waits until all operations associated with the `async-argument`, that is, stream, have completed.

Normally, the management of the streams that are created as a result of using the `async` clause, is done without any intervention by the caller. This implies the association between the `async-argument` and the CUDA stream is maintained for the lifetime of the program. However, this association can be changed through the use of the library function `acc_set_cuda_stream`. When the function `acc_set_cuda_stream` is called, the CUDA stream that was originally associated with the `async` clause is destroyed. Caution should be taken when changing the association as subsequent references to the `async-argument` refer to a different CUDA stream.

¹ See "Stream Management" in "CUDA Driver API", TRM-06703-001, Version 5.5, for additional information

9 OpenACC Library Interoperability

9.1 Introduction

The OpenACC library uses the CUDA Driver API, and may interact with programs that use the Runtime library directly, or another library based on the Runtime library, e.g., CUBLAS¹. This chapter describes the use cases and what changes are required in order to use both the OpenACC library and the CUBLAS and Runtime libraries within a program.

9.2 First invocation: NVIDIA CUBLAS library API

In this first use case (see below), a function in the CUBLAS library is called prior to any of the functions in the OpenACC library. More specifically, the function `cublasCreate()`.

When invoked, the function initializes the library and allocates the hardware resources on the host and the device on behalf of the caller. Once the initialization and allocation has completed, a handle is returned to the caller. The OpenACC library also requires initialization and allocation of hardware resources. Since the CUBLAS library has already allocated the hardware resources for the device, all that is left to do is to initialize the OpenACC library and acquire the hardware resources on the host.

Prior to calling the OpenACC function that initializes the library and allocate the host hardware resources, you need to acquire the device number that was allocated during the call to `cublasCreate()`. The invoking of the runtime library function `cudaGetDevice()` accomplishes this. Once acquired, the device number is passed along with the device type as parameters to the OpenACC library function `acc_set_device_num()`.

Once the call to `acc_set_device_num()` has completed, the OpenACC library uses the context that was created during the call to `cublasCreate()`. In other words, both libraries share the same context.

```
/* Create the handle */
s = cublasCreate(&h);
if (s != CUBLAS_STATUS_SUCCESS)
{
    fprintf(stderr, "cublasCreate failed %d\n", s);
    exit(EXIT_FAILURE);
}

/* Get the device number */
e = cudaGetDevice(&dev);
if (e != cudaSuccess)
{
    fprintf(stderr, "cudaGetDevice failed %d\n", e);
    exit(EXIT_FAILURE);
}

/* Initialize OpenACC library and use device 'dev' */
acc_set_device_num(dev, acc_device_nvidia);
```

Use Case 1

¹ See section 2.26, "Interactions with the CUDA Driver API" in "CUDA Runtime API", Version 5.5, and section 2.27, "VDPAU Interoperability", in "CUDA Driver API", TRM-06703-001, Version 5.5, for additional information on library interoperability.

9.3 First invocation: OpenACC library API

In this second use case (see below), a function in the OpenACC library is called prior to any of the functions in the CUBLAS library. More specifically, the function `acc_set_device_num()`.

In the use case presented here, the function `acc_set_device_num()` is used to both initialize the OpenACC library and allocate the hardware resources on the host and the device. In the call to the function, the call parameters specify which device to use and what device type to use, i.e., `acc_device_nvidia`. It should be noted that this is but one method to initialize the OpenACC library and allocate the appropriate hardware resources. Other methods are available through the use of environment variables and these is discussed in the next section.

Once the call to `acc_set_device_num()` has completed, other OpenACC functions can be called as seen with multiple calls being made to `acc_copyin()`. In addition, calls can be made to functions in the CUBLAS library. In the use case a call to `cublasCreate()` is made subsequent to the calls to `acc_copyin()`. As seen in the previous use case, a call to `cublasCreate()` initializes the CUBLAS library and allocates the hardware resources on the host and the device. However, since the device has already been allocated, `cublasCreate()` only initializes the CUBLAS library and allocates the appropriate hardware resources on the host. The context that was created as part of the OpenACC initialization is shared with the CUBLAS library, similarly to the first use case.

```
dev = 0;

acc_set_device_num(dev, acc_device_nvidia);

/* Copy the first set to the device */
d_X = acc_copyin(&h_X[0], N * sizeof (float));
if (d_X == NULL)
{
    fprintf(stderr, "copyin error h_X\n");
    exit(EXIT_FAILURE);
}

/* Copy the second set to the device */
d_Y = acc_copyin(&h_Y1[0], N * sizeof (float));
if (d_Y == NULL)
{
    fprintf(stderr, "copyin error h_Y1\n");
    exit(EXIT_FAILURE);
}

/* Create the handle */
s = cublasCreate(&h);
if (s != CUBLAS_STATUS_SUCCESS)
{
    fprintf(stderr, "cublasCreate failed %d\n", s);
    exit(EXIT_FAILURE);
}

/* Perform saxpy using CUBLAS library function */
s = cublasSaxpy(h, N, &alpha, d_X, 1, d_Y, 1);
if (s != CUBLAS_STATUS_SUCCESS)
{
```

```
        fprintf(stderr, "cublasSaxpy failed %d\n", s);
        exit(EXIT_FAILURE);
    }

    /* Copy the results from the device */
    acc_memcpy_from_device(&h_Y1[0], d_Y, N * sizeof (float));
```

Use Case 2

9.4 OpenACC library and environment variables

There are two environment variables associated with the OpenACC library that may be used to control the device type and device number: `ACC_DEVICE_TYPE` and `ACC_DEVICE_NUM`, respectively. These two environment variables can be used as an alternative to calling `acc_set_device_num()`. As seen in the second use case, the device type and device number were specified using `acc_set_device_num()`. If however, the aforementioned environment variables were set, then the call to `acc_set_device_num()` would not be required.

The use of the environment variables is only relevant when an OpenACC function is called prior to a call to `cudaCreate()`. If `cudaCreate()` is called prior to a call to an OpenACC function, then you must call `acc_set_device_num()`²

² More complete information about `ACC_DEVICE_TYPE` and `ACC_DEVICE_NUM` can be found in sections 4.1 and 4.2 of the OpenACC (<https://www.openacc.org>) Application Programming Interface”, Version 2.6.

10 OpenACC Profiling Interface

10.1 Implementation Status and Implementation-Defined Behavior

We’re implementing the OpenACC Profiling Interface as defined by the OpenACC 2.6 specification. We’re clarifying some aspects here as *implementation-defined behavior*, while they’re still under discussion within the OpenACC Technical Committee.

This implementation is tuned to keep the performance impact as low as possible for the (very common) case that the Profiling Interface is not enabled. This is relevant, as the Profiling Interface affects all the *hot* code paths (in the target code, not in the offloaded code). Users of the OpenACC Profiling Interface can be expected to understand that performance is impacted to some degree once the Profiling Interface is enabled: for example, because of the *runtime* (libgomp) calling into a third-party *library* for every event that has been registered.

We’re not yet accounting for the fact that *OpenACC events may occur during event processing*. We just handle one case specially, as required by CUDA 9.0 `nvprof`, that `acc_get_device_type` (Section 6.3 [`acc_get_device_type`], page 73)) may be called from `acc_ev_device_init_start`, `acc_ev_device_init_end` callbacks.

We’re not yet implementing initialization via a `acc_register_library` function that is either statically linked in, or dynamically via `LD_PRELOAD`. Initialization via `acc_register_library` functions dynamically loaded via the `ACC_PROFLIB` environment variable does work, as does directly calling `acc_prof_register`, `acc_prof_unregister`, `acc_prof_lookup`.

As currently there are no inquiry functions defined, calls to `acc_prof_lookup` always returns `NULL`.

There aren’t separate *start*, *stop* events defined for the event types `acc_ev_create`, `acc_ev_delete`, `acc_ev_alloc`, `acc_ev_free`. It’s not clear if these should be triggered before or after the actual device-specific call is made. We trigger them after.

Remarks about data provided to callbacks:

`acc_prof_info.event_type`

It’s not clear if for *nested* event callbacks (for example, `acc_ev_enqueue_launch_start` as part of a parent compute construct), this should be set for the nested event (`acc_ev_enqueue_launch_start`), or if the value of the parent construct should remain (`acc_ev_compute_construct_start`). In this implementation, the value generally corresponds to the innermost nested event type.

`acc_prof_info.device_type`

- For `acc_ev_compute_construct_start`, and in presence of an `if` clause with *false* argument, this still refers to the offloading device type. It’s not clear if that’s the expected behavior.
- Complementary to the item before, for `acc_ev_compute_construct_end`, this is set to `acc_device_host` in presence of an `if` clause with *false* argument. It’s not clear if that’s the expected behavior.

`acc_prof_info.thread_id`

Always -1; not yet implemented.

`acc_prof_info.async`

- Not yet implemented correctly for `acc_ev_compute_construct_start`.
- In a compute construct, for host-fallback execution/`acc_device_host` it always is `acc_async_sync`. It is unclear if that is the expected behavior.
- For `acc_ev_device_init_start` and `acc_ev_device_init_end`, it will always be `acc_async_sync`. It is unclear if that is the expected behavior.

`acc_prof_info.async_queue`

There is no *limited number of asynchronous queues* in libgomp. This always has the same value as `acc_prof_info.async`.

`acc_prof_info.src_file`

Always NULL; not yet implemented.

`acc_prof_info.func_name`

Always NULL; not yet implemented.

`acc_prof_info.line_no`

Always -1; not yet implemented.

`acc_prof_info.end_line_no`

Always -1; not yet implemented.

`acc_prof_info.func_line_no`

Always -1; not yet implemented.

`acc_prof_info.func_end_line_no`

Always -1; not yet implemented.

`acc_event_info.event_type`, `acc_event_info.*.event_type`

Relating to `acc_prof_info.event_type` discussed above, in this implementation, this will always be the same value as `acc_prof_info.event_type`.

`acc_event_info.*.parent_construct`

- Will be `acc_construct_parallel` for all OpenACC compute constructs as well as many OpenACC Runtime API calls; should be the one matching the actual construct, or `acc_construct_runtime_api`, respectively.
- Will be `acc_construct_enter_data` or `acc_construct_exit_data` when processing variable mappings specified in OpenACC *declare* directives; should be `acc_construct_declare`.
- For implicit `acc_ev_device_init_start`, `acc_ev_device_init_end`, and explicit as well as implicit `acc_ev_alloc`, `acc_ev_free`, `acc_ev_enqueue_upload_start`, `acc_ev_enqueue_upload_end`, `acc_ev_enqueue_download_start`, and `acc_ev_enqueue_download_end`, will be `acc_construct_parallel`; should reflect the real parent construct.

`acc_event_info.*.implicit`

For `acc_ev_alloc`, `acc_ev_free`, `acc_ev_enqueue_upload_start`, `acc_ev_enqueue_upload_end`, `acc_ev_enqueue_download_start`, and

`acc_ev_enqueue_download_end`, this currently will be 1 also for explicit usage.

`acc_event_info.data_event.var_name`

Always NULL; not yet implemented.

`acc_event_info.data_event.host_ptr`

For `acc_ev_alloc`, and `acc_ev_free`, this is always NULL.

`typedef union acc_api_info`

... as printed in 5.2.3. *Third Argument: API-Specific Information*. This should obviously be `typedef struct acc_api_info`.

`acc_api_info.device_api`

Possibly not yet implemented correctly for `acc_ev_compute_construct_start`, `acc_ev_device_init_start`, `acc_ev_device_init_end`: will always be `acc_device_api_none` for these event types. For `acc_ev_enter_data_start`, it will be `acc_device_api_none` in some cases.

`acc_api_info.device_type`

Always the same as `acc_prof_info.device_type`.

`acc_api_info.vendor`

Always -1; not yet implemented.

`acc_api_info.device_handle`

Always NULL; not yet implemented.

`acc_api_info.context_handle`

Always NULL; not yet implemented.

`acc_api_info.async_handle`

Always NULL; not yet implemented.

Remarks about certain event types:

`acc_ev_device_init_start`, `acc_ev_device_init_end`

- When a compute construct triggers implicit `acc_ev_device_init_start` and `acc_ev_device_init_end` events, they currently aren't *nested within* the corresponding `acc_ev_compute_construct_start` and `acc_ev_compute_construct_end`, but they're currently observed *before* `acc_ev_compute_construct_start`. It's not clear what to do: the standard asks us provide a lot of details to the `acc_ev_compute_construct_start` callback, without (implicitly) initializing a device before?
- Callbacks for these event types will not be invoked for calls to the `acc_set_device_type` and `acc_set_device_num` functions. It's not clear if they should be.

`acc_ev_enter_data_start`, `acc_ev_enter_data_end`, `acc_ev_exit_data_start`, `acc_ev_exit_data_end`

- Callbacks for these event types will also be invoked for OpenACC *host_data* constructs. It's not clear if they should be.

- Callbacks for these event types will also be invoked when processing variable mappings specified in OpenACC *declare* directives. It's not clear if they should be.

Callbacks for the following event types will be invoked, but dispatch and information provided therein has not yet been thoroughly reviewed:

- `acc_ev_alloc`
- `acc_ev_free`
- `acc_ev_update_start`, `acc_ev_update_end`
- `acc_ev_enqueue_upload_start`, `acc_ev_enqueue_upload_end`
- `acc_ev_enqueue_download_start`, `acc_ev_enqueue_download_end`

During device initialization, and finalization, respectively, callbacks for the following event types will not yet be invoked:

- `acc_ev_alloc`
- `acc_ev_free`

Callbacks for the following event types have not yet been implemented, so currently won't be invoked:

- `acc_ev_device_shutdown_start`, `acc_ev_device_shutdown_end`
- `acc_ev_runtime_shutdown`
- `acc_ev_create`, `acc_ev_delete`
- `acc_ev_wait_start`, `acc_ev_wait_end`

For the following runtime library functions, not all expected callbacks will be invoked (mostly concerning implicit device initialization):

- `acc_get_num_devices`
- `acc_set_device_type`
- `acc_get_device_type`
- `acc_set_device_num`
- `acc_get_device_num`
- `acc_init`
- `acc_shutdown`

Aside from implicit device initialization, for the following runtime library functions, no callbacks will be invoked for shared-memory offloading devices (it's not clear if they should be):

- `acc_malloc`
- `acc_free`
- `acc_copyin`, `acc_present_or_copyin`, `acc_copyin_async`
- `acc_create`, `acc_present_or_create`, `acc_create_async`
- `acc_copyout`, `acc_copyout_async`, `acc_copyout_finalize`, `acc_copyout_finalize_async`
- `acc_delete`, `acc_delete_async`, `acc_delete_finalize`, `acc_delete_finalize_async`

- `acc_update_device`, `acc_update_device_async`
- `acc_update_self`, `acc_update_self_async`
- `acc_map_data`, `acc_unmap_data`
- `acc_memcpy_to_device`, `acc_memcpy_to_device_async`
- `acc_memcpy_from_device`, `acc_memcpy_from_device_async`

11 OpenMP-Implementation Specifics

11.1 Implementation-defined ICV Initialization

<i>affinity-format-var</i>	See Section 4.2 [OMP_AFFINITY_FORMAT], page 60.
<i>def-allocator-var</i>	See Section 4.1 [OMP_ALLOCATOR], page 59.
<i>max-active-levels-var</i>	See Section 4.8 [OMP_MAX_ACTIVE_LEVELS], page 62.
<i>dyn-var</i>	See Section 4.7 [OMP_DYNAMIC], page 62.
<i>nthreads-var</i>	See Section 4.12 [OMP_NUM_THREADS], page 63.
<i>num-devices-var</i>	Number of non-host devices found by GCC's run-time library
<i>num-procs-var</i>	The number of CPU cores on the initial device, except that affinity settings might lead to a smaller number. On non-host devices, the value of the <i>nthreads-var</i> ICV.
<i>place-partition-var</i>	See Section 4.14 [OMP_PLACES], page 64.
<i>run-sched-var</i>	See Section 4.16 [OMP_SCHEDULE], page 65.
<i>stacksize-var</i>	See Section 4.15 [OMP_STACKSIZE], page 65.
<i>thread-limit-var</i>	See Section 4.18 [OMP_TEAMS_THREAD_LIMIT], page 66.
<i>wait-policy-var</i>	See Section 4.20 [OMP_WAIT_POLICY], page 67, and Section 4.24 [GOMP_SPINCOUNT], page 68,

11.2 OpenMP Context Selectors

vendor is always **gnu**. References are to the GCC manual.

For the host compiler, **kind** always matches **host**, **cpu** and **any**; for the offloading architectures AMD GCN and Nvidia PTX, **kind** always matches **nohost**, **gpu** and **any**. For the x86 family of computers, AMD GCN and Nvidia PTX the following traits are supported in addition; while OpenMP is supported on more architectures, GCC currently does not match any **arch** or **isa** traits for those.

Note that for AMD GCN and Nvidia PTX, the **isa** is currently an exact match between the compiled-for ISA architecture and the matching **isa** trait value. For instance, when compiling for **gfx942**, the **isa** trait value **gfx9-4-generic** is not matched and, likewise, **gfx942** is not matched when compiling for its generic architecture.

arch	isa
x86 , x86_64 , i386 , i486 , i586 , i686 , ia32	See -m... flags in “x86 Options” (without -m)
amdgc n, gc n	See -march= in “AMD GCN Options”
nvpt x, nvptx64	See -march= in “Nvidia PTX Options”

11.3 Memory allocation

The description below applies to:

- Explicit use of the OpenMP API routines, see Section 3.12 [Memory Management Routines], page 50.

- The `allocate` clause, except when the `allocator` modifier is a constant expression with value `omp_default_mem_alloc` and no `align` modifier has been specified. (In that case, the normal `malloc` allocation is used.)
- The `allocate` directive for variables in static memory; while the alignment is honored, the normal static memory is used.
- Using the `allocate` directive for automatic/stack variables, except when the `allocator` clause is a constant expression with value `omp_default_mem_alloc` and no `align` clause has been specified. (In that case, the normal allocation is used: stack allocation and, sometimes for Fortran, also `malloc` [depending on flags such as `-fstack-arrays`].)
- In Fortran, the `allocators` directive and the executable `allocate` directive for Fortran pointers and allocatables is supported, but requires that files containing those directives has to be compiled with `-fopenmp-allocators`. Additionally, all files that might explicitly or implicitly deallocate memory allocated that way must also be compiled with that option.
- The used alignment is the maximum of the value the `align` clause and the alignment of the type after honoring, if present, the `aligned` (GNU::`aligned`) attribute and C's `_Alignas` and C++'s `alignas`. However, the `align` clause of the `allocate` directive has no effect on the value of C's `_Alignof` and C++'s `alignof`.

GCC supports the following predefined allocators and predefined memory spaces:

Predefined allocators	Associated predefined memory spaces
<code>omp_default_mem_alloc</code>	<code>omp_default_mem_space</code>
<code>omp_large_cap_mem_alloc</code>	<code>omp_large_cap_mem_space</code>
<code>omp_const_mem_alloc</code>	<code>omp_const_mem_space</code>
<code>omp_high_bw_mem_alloc</code>	<code>omp_high_bw_mem_space</code>
<code>omp_low_lat_mem_alloc</code>	<code>omp_low_lat_mem_space</code>
<code>omp_cgroup_mem_alloc</code>	<code>omp_low_lat_mem_space</code> (implementation defined)
<code>omp_pteam_mem_alloc</code>	<code>omp_low_lat_mem_space</code> (implementation defined)
<code>omp_thread_mem_alloc</code>	<code>omp_low_lat_mem_space</code> (implementation defined)
<code>ompx_gnu_pinned_mem_alloc</code>	<code>omp_default_mem_space</code> (GNU extension)
<code>ompx_gnu_managed_mem_alloc</code>	<code>ompx_gnu_managed_mem_space</code> (GNU extension)

Each predefined allocator, including `omp_null_allocator`, has a corresponding allocator class template that meet the C++ allocator completeness requirements. These are located in the `omp::allocator` namespace, and the `ompx::allocator` namespace for gnu extensions. This allows the allocator-aware C++ standard library containers to use OpenMP allocation routines; for instance:

```
std::vector<int, omp::allocator::cgroup_mem<int>> vec;
```

The following allocator templates are supported:

Predefined allocators	Associated allocator template
<code>omp_null_allocator</code>	<code>omp::allocator::null_allocator</code>

<code>omp_default_mem_alloc</code>	<code>omp::allocator::default_mem</code>
<code>omp_large_cap_mem_alloc</code>	<code>omp::allocator::large_cap_mem</code>
<code>omp_const_mem_alloc</code>	<code>omp::allocator::const_mem</code>
<code>omp_high_bw_mem_alloc</code>	<code>omp::allocator::high_bw_mem</code>
<code>omp_low_lat_mem_alloc</code>	<code>omp::allocator::low_lat_mem</code>
<code>omp_cgroup_mem_alloc</code>	<code>omp::allocator::cgroup_mem</code>
<code>omp_pteam_mem_alloc</code>	<code>omp::allocator::pteam_mem</code>
<code>omp_thread_mem_alloc</code>	<code>omp::allocator::thread_mem</code>
<code>ompx_gnu_pinned_mem_alloc</code>	<code>ompx::allocator::gnu_pinned_mem</code>
<code>ompx_gnu_managed_mem_alloc</code>	<code>ompx::allocator::gnu_managed_mem</code>

The following traits are available when constructing a new allocator; if a trait is not specified or with the value `default`, the specified default value is used for that trait. The predefined allocators use the default values of each trait, except that the `omp_cgroup_mem_alloc`, `omp_pteam_mem_alloc`, and `omp_thread_mem_alloc` allocators have the `access` trait set to `cgroup`, `pteam`, and `thread`, respectively. For each trait, a named constant prefixed by `omp_atk_` exists; for each non-numeric value, a named constant prefixed by `omp_atv_` exists.

Trait	Allowed values	Default value
<code>sync_hint</code>	<code>contended</code> , <code>uncontended</code> , <code>serialized</code> , <code>private</code>	<code>contended</code>
<code>alignment</code>	Positive integer being a power of two	1 byte
<code>access</code>	<code>all</code> , <code>cgroup</code> , <code>pteam</code> , <code>thread</code>	<code>all</code>
<code>pool_size</code>	Positive integer (bytes)	See below.
<code>fallback</code>	<code>default_mem_fb</code> , <code>null_fb</code> , <code>abort_fb</code> , <code>allocator_fb</code>	See below
<code>fb_data</code>	<i>allocator handle</i>	(none)
<code>pinned</code>	<code>true</code> , <code>false</code>	See below
<code>partition</code>	<code>environment</code> , <code>nearest</code> , <code>blocked</code> , <code>interleaved</code>	<code>environment</code>

For the `fallback` trait, the default value is `null_fb` for the `omp_default_mem_alloc` allocator and any allocator that is associated with device memory; for all other allocators, it is `default_mem_fb` by default.

For the `pinned` trait, the default value is `true` for predefined allocator `ompx_gnu_pinned_mem_alloc` (a GNU extension), and `false` for all others.

The following description applies to the initial device (the host) and largely also to non-host devices; for the latter, also see Chapter 12 [Offload-Target Specifics], page 113.

For the memory spaces, the following applies:

- `omp_default_mem_space` is supported
- `omp_const_mem_space` maps to `omp_default_mem_space`
- `omp_low_lat_mem_space` is only available on supported devices, and maps to `omp_default_mem_space` otherwise.
- `omp_large_cap_mem_space` maps to `omp_default_mem_space`, unless the `memkind` library is available
- `omp_high_bw_mem_space` maps to `omp_default_mem_space`, unless the `memkind` library is available

- `ompx_gnu_managed_mem_space` is a GNU extension that provides managed memory accessible by both host and devices. The memory space is available if the offload target associated with the *default-device-var* ICV supports managed memory (see Chapter 12 [Offload-Target Specifics], page 113). This memory is accessible by both the host and the device at the same address, so it need not be mapped with `map` clauses. Instead, use the `is_device_ptr` clause or `has_device_addr` clause to indicate that the pointer is already accessible on the device. If managed memory is not supported by the default device, as configured at the moment the allocator is called, then the allocator will use the fall-back setting. If the default device is configured differently when the memory is freed, via `omp_free` or `omp_realloc`, the result may be undefined.

On Linux systems, where the memkind library (<https://github.com/memkind/memkind>) (`libmemkind.so.0`) is available at runtime and the respective memkind kind is supported, it is used when creating memory allocators requesting

- the `partition` trait `interleaved` except when the memory space is `omp_large_cap_mem_space` (uses `MEMKIND_HBW_INTERLEAVE`)
- the memory space is `omp_high_bw_mem_space` (uses `MEMKIND_HBW_PREFERRED`)
- the memory space is `omp_large_cap_mem_space` (uses `MEMKIND_DAX_KMEM_ALL` or, if not available, `MEMKIND_DAX_KMEM`)

On Linux systems, where the numa library (<https://github.com/numactl/numactl>) (`libnuma.so.1`) is available at runtime, it is used when creating memory allocators requesting

- the `partition` trait `nearest`, except when both the libmemkind library is available and the memory space is either `omp_large_cap_mem_space` or `omp_high_bw_mem_space`

Note that the numa library will round up the allocation size to a multiple of the system page size; therefore, consider using it only with large data or by sharing allocations via the `pool_size` trait. Furthermore, the Linux kernel does not guarantee that an allocation will always be on the nearest NUMA node nor that after reallocation the same node will be used. Note additionally that, on Linux, the default setting of the memory placement policy is to use the current node; therefore, unless the memory placement policy has been overridden, the `partition` trait `environment` (the default) will be effectively a `nearest` allocation.

Additional notes regarding the traits:

- The `pinned` trait is supported on Linux hosts, but is usually subject to the OS `ulimit/rlimit` locked memory settings (see Chapter 12 [Offload-Target Specifics], page 113, for exceptions). The implementation uses a custom allocator to try to use as few memory pages as possible. At present, freed pinned memory is not returned to the OS (although it may be reused by subsequent pinned allocations).
- The default for the `pool_size` trait is no pool and for every (re)allocation the associated library routine is called, which might internally use a memory pool. Currently, the same applies when a `pool_size` has been specified, except that once allocations exceed the pool size, the action of the `fallback` trait applies.
- For the `partition` trait, the partition part size will be the same as the requested size (i.e. `interleaved` or `blocked` has no effect), except for `interleaved` when the memkind library is available. Furthermore, for `nearest` and unless the numa library is available, the memory might not be on the same NUMA node as thread that allocated

the memory; on Linux, this is in particular the case when the memory placement policy is set to preferred.

- The `access` trait has no effect such that memory is always accessible by all threads. (Except on supported no-host devices.)
- The `sync_hint` trait has no effect.

See also: Chapter 12 [Offload-Target Specifics], page 113,

12 Offload-Target Specifics

The following sections present notes on the offload-target specifics

12.1 AMD Radeon (GCN)

On the hardware side, there is the hierarchy (fine to coarse):

- work item (thread)
- wavefront
- work group
- compute unit (CU)

All OpenMP and OpenACC levels are used, i.e.

- OpenMP's `simd` and OpenACC's `vector` map to work items (thread)
- OpenMP's threads ("parallel") and OpenACC's workers map to wavefronts
- OpenMP's teams and OpenACC's gang use a threadpool with the size of the number of teams or gangs, respectively.

The used sizes are

- Number of teams is the specified `num_teams` (OpenMP) or `num_gangs` (OpenACC) or otherwise the number of CU. It is limited by two times the number of CU.
- Number of wavefronts is 4 for gfx900 and 16 otherwise; `num_threads` (OpenMP) and `num_workers` (OpenACC) overrides this if smaller.
- The wavefront has 102 scalars and 64 vectors
- Number of workitems is always 64
- The hardware permits maximally 40 workgroups/CU and 16 wavefronts/workgroup up to a limit of 40 wavefronts in total per CU.
- 80 scalars registers and 24 vector registers in non-kernel functions (the chosen procedure-calling API).
- For the kernel itself: as many as register pressure demands (number of teams and number of threads, scaled down if registers are exhausted)

The implementation remark:

- I/O within OpenMP target regions and OpenACC compute regions is supported using the C library `printf` functions and the Fortran `print/write` statements.
- Reverse offload regions (i.e. `target` regions with `device(ancestor:1)`) are processed serially per `target` region such that the next reverse offload region is only executed after the previous one returned.
- OpenMP code that has a `requires` directive with `self_maps` or `unified_shared_memory` is only supported if all AMD GPUs have the `HSA_AMD_SYSTEM_INFO_SVM_ACCESSIBLE_BY_DEFAULT` property; for discrete GPUs, this may require setting the `HSA_XNACK` environment variable to '1'; for systems with both an APU and a discrete GPU that does not support XNACK, consider using `ROCR_VISIBLE_DEVICES` to enable only the APU. If not supported, all AMD GPU devices are removed from the list of available devices ("host fallback").

- The available stack size can be changed using the `GCN_STACK_SIZE` environment variable; the default is 32 kiB per thread.
- Low-latency memory (`omp_low_lat_mem_space`) is supported when the `access` trait is set to `cgroup`. The default pool size is automatically scaled to share the 64 kiB LDS memory between the number of teams configured to run on each compute-unit, but may be adjusted at runtime by setting environment variable `GOMP_GCN_LOWLAT_POOL=bytes`.
- `omp_low_lat_mem_alloc` cannot be used with true low-latency memory because the definition implies the `omp_atv_all` trait; main graphics memory is used instead.
- `omp_cgroup_mem_alloc`, `omp_pteam_mem_alloc`, and `omp_thread_mem_alloc`, all use low-latency memory as first preference, and fall back to main graphics memory when the low-latency pool is exhausted.
- Pinned memory allocated using `omp_alloc` with the `ompx_gnu_pinned_mem_alloc` allocator or the `pinned` trait is obtained via the CUDA API when an NVPTX device is present. This provides a performance boost for NVPTX offload code and also allows unlimited use of pinned memory regardless of the OS `ulimit/rlimit` settings.
- Managed memory allocated with the OpenMP `ompx_gnu_managed_mem_alloc` allocator or in the `ompx_gnu_managed_mem_space` is not currently supported for AMD GPU devices; attempting to use it in an allocator will trigger the fall-back trait.
- The OpenMP routines `omp_target_memcpy_rect` and `omp_target_memcpy_rect_async` and the `target update` directive for non-contiguous list items use the 3D memory-copy function of the HSA library. Higher dimensions call this functions in a loop and are therefore supported.
- The unique identifier (UID), used with OpenMP's API UID routines, is the value returned by the HSA runtime library for `HSA_AMD_AGENT_INFO_UUID`. For GPUs, it is currently 'GPU-' followed by 16 lower-case hex digits, yielding a string like GPU-f914a2142fc3413a. The output matches the one used by `rocminfo`.

12.1.1 OpenMP interop – Foreign-Runtime Support for AMD GPUs

On AMD GPUs, the foreign runtimes are HIP (C++ Heterogeneous-Compute Interface for Portability) and HSA (Heterogeneous System Architecture), where HIP is the default. The interop object is created using OpenMP's `interop` directive or, implicitly, when invoking a `declare variant` procedure that has the `append_args` clause. In either case, the `prefer_type` modifier determines whether HIP or HSA is used.

When specifying the `targetsync` modifier: For HIP, a stream is created using `hipStreamCreate`. For HSA, a queue is created of type `HSA_QUEUE_TYPE_MULTI` with a queue size of 64.

Invoke the Section 3.11 [Interoperability Routines], page 46, on an interop object to obtain the following properties. For properties with integral (int), pointer (ptr), or string (str) data type, call `omp_get_interop_int`, `omp_get_interop_ptr`, or `omp_get_interop_str`, respectively. Note that `device_num` is the OpenMP device number while `device` is the HIP device number or HSA device handle.

When using HIP with C and C++, the `__HIP_PLATFORM_AMD__` preprocessor macro must be defined before including the HIP header files.

For the API routine call, add the prefix `omp_ipr_` to the property name; for instance:

```
omp_interop_rc_t ret;
int device_num = omp_get_interop_int (my_interop_obj, omp_ipr_device_num, &ret);
```

Available properties for an HIP interop object:

Property	C data type	API routine	value constant)	(if
<code>fr_id</code>	<code>omp_interop_fr_t</code>	<code>int</code>	<code>omp_fr_hip</code>	
<code>fr_name</code>	<code>const char *</code>	<code>str</code>	<code>"hip"</code>	
<code>vendor</code>	<code>int</code>	<code>int</code>	<code>1</code>	
<code>vendor_name</code>	<code>const char *</code>	<code>str</code>	<code>"amd"</code>	
<code>device_num</code>	<code>int</code>	<code>int</code>		
<code>platform</code>	<code>N/A</code>			
<code>device</code>	<code>hipDevice_t</code>	<code>int</code>		
<code>device_context</code>	<code>hipCtx_t</code>	<code>ptr</code>		
<code>targetsync</code>	<code>hipStream_t</code>	<code>ptr</code>		

Available properties for an HSA interop object:

Property	C data type	API routine	value constant)	(if
<code>fr_id</code>	<code>omp_interop_fr_t</code>	<code>int</code>	<code>omp_fr_hsa</code>	
<code>fr_name</code>	<code>const char *</code>	<code>str</code>	<code>"hsa"</code>	
<code>vendor</code>	<code>int</code>	<code>int</code>	<code>1</code>	
<code>vendor_name</code>	<code>const char *</code>	<code>str</code>	<code>"amd"</code>	
<code>device_num</code>	<code>int</code>	<code>int</code>		
<code>platform</code>	<code>N/A</code>			
<code>device</code>	<code>hsa_agent *</code>	<code>ptr</code>		
<code>device_context</code>	<code>N/A</code>			
<code>targetsync</code>	<code>hsa_queue *</code>	<code>ptr</code>		

12.2 nvptx

On the hardware side, there is the hierarchy (fine to coarse):

- thread
- warp
- thread block
- streaming multiprocessor

All OpenMP and OpenACC levels are used, i.e.

- OpenMP's `simd` and OpenACC's vector map to threads
- OpenMP's threads ("parallel") and OpenACC's workers map to warps
- OpenMP's teams and OpenACC's gang use a threadpool with the size of the number of teams or gangs, respectively.

The used sizes are

- The `warp_size` is always 32
- CUDA kernel launched: `dim={#teams,1,1}`, `blocks={#threads,warp_size,1}`.

- The number of teams is limited by the number of blocks the device can host simultaneously.

Additional information can be obtained by setting the environment variable to `GOMP_DEBUG=1` (very verbose; grep for `kernel.*launch` for launch parameters).

GCC generates generic PTX ISA code, which is just-in-time compiled by CUDA, which caches the JIT in the user's directory (see CUDA documentation; can be tuned by the environment variables `CUDA_CACHE_{DISABLE,MAXSIZE,PATH}`).

Note: While PTX ISA is generic, the `-mptx=` and `-march=` commandline options still affect the used PTX ISA code and, thus, the requirements on CUDA version and hardware.

The implementation remark:

- I/O within OpenMP target regions and OpenACC compute regions is supported using the C library `printf` functions. Additionally, the Fortran `print/write` statements are supported within OpenMP target regions, but not yet within OpenACC compute regions.
- Compilation OpenMP code that contains `requires reverse_offload` requires at least `-march=sm_35`, compiling for `-march=sm_30` is not supported.
- For code containing reverse offload (i.e. `target` regions with `device(ancestor:1)`), there is a slight performance penalty for *all* target regions, consisting mostly of shut-down delay Per device, reverse offload regions are processed serially such that the next reverse offload region is only executed after the previous one returned.
- OpenMP code that has a `requires` directive with `self_maps` or `unified_shared_memory` runs on nvptx devices if and only if all of those support the `pageableMemoryAccess` property;¹ otherwise, all nvptx device are removed from the list of available devices (“host fallback”).
- The default per-warp stack size is 128 kiB; see also `-msoft-stack` in the GCC manual.
- Low-latency memory (`omp_low_lat_mem_space`) is supported when the `access` trait is set to `cgroup`, and libgomp has been built for PTX ISA version 4.1 or higher (such as in GCC's default configuration). The default pool size is 8 kiB per team, but may be adjusted at runtime by setting environment variable `GOMP_NVPTX_LOWLAT_POOL=bytes`. The maximum value is limited by the available hardware, and care should be taken that the selected pool size does not unduly limit the number of teams that can run simultaneously.
- `omp_low_lat_mem_alloc` cannot be used with true low-latency memory because the definition implies the `omp_atv_all` trait; main graphics memory is used instead.
- `omp_cgroup_mem_alloc`, `omp_pteam_mem_alloc`, and `omp_thread_mem_alloc`, all use low-latency memory as first preference, and fall back to main graphics memory when the low-latency pool is exhausted.
- Managed memory allocated on the host with the `ompx_gnu_managed_mem_alloc` allocator or in the `ompx_gnu_managed_mem_space` (both GNU extensions) allocate memory in the CUDA Managed Memory space using `cuMemAllocManaged`. This memory is accessible by both the host and the device at the same address, so it need not be mapped with `map` clauses. Instead, use the `is_device_ptr` clause or `has_device_addr` clause

¹ <https://docs.nvidia.com/cuda/cuda-c-programming-guide/index.html#um-requirements>

to indicate that the pointer is already accessible on the device. The CUDA runtime will automatically handle data migration between host and device as needed. If managed memory is not supported by the default device, as configured at the moment the allocator is called, then the allocator will use the fall-back setting. If the default device is configured differently when the memory is freed, via `omp_free` or `omp_realloc`, the result may be undefined.

- The OpenMP routines `omp_target_memcpy_rect` and `omp_target_memcpy_rect_async` and the `target update` directive for non-contiguous list items use the 2D and 3D memory-copy functions of the CUDA library. Higher dimensions call those functions in a loop and are therefore supported.
- The unique identifier (UID), used with OpenMP's API UID routines, consists of the 'GPU-' prefix followed by the 16-bytes UUID as returned by the CUDA runtime library. This UUID is output in grouped lower-case hex digits; the grouping of those 32 digits is: 8 digits, hyphen, 4 digits, hyphen, 4 digits, hyphen, 16 digits. This leads to a string like `GPU-a8081c9e-f03e-18eb-1827-bf5ba95afa5d`. The output matches the format used by `nvidia-smi`.

12.2.1 OpenMP interop – Foreign-Runtime Support for Nvidia GPUs

On Nvidia GPUs, the foreign runtimes APIs are the CUDA runtime API, the CUDA driver API, and HIP, the C++ Heterogeneous-Compute Interface for Portability that is—on CUDA-based systems—a very thin layer on top of the CUDA API. By default, CUDA is used. The interop object is created using OpenMP's `interop` directive or, implicitly, when invoking a `declare variant` procedure that has the `append_args` clause. In either case, the `prefer_type` modifier determines whether CUDA, CUDA driver, or HSA is used.

When specifying the `targetsync` modifier, a CUDA stream is created using the `CU_STREAM_DEFAULT` flag.

Invoke the Section 3.11 [Interoperability Routines], page 46, on an interop object to obtain the following properties. For properties with integral (int), pointer (ptr), or string (str) data type, call `omp_get_interop_int`, `omp_get_interop_ptr`, or `omp_get_interop_str`, respectively. Note that `device_num` is the OpenMP device number while `device` is the CUDA, CUDA Driver, or HIP device number.

When using HIP with C and C++, the `__HIP_PLATFORM_NVIDIA__` preprocessor macro must be defined before including the HIP header files.

For the API routine call, add the prefix `omp_ipr_` to the property name; for instance:

```
omp_interop_rc_t ret;
int device_num = omp_get_interop_int (my_interop_obj, omp_ipr_device_num, &ret);
```

Available properties for a CUDA runtime API interop object:

Property	C data type	API routine	value (if constant)
<code>fr_id</code>	<code>omp_interop_fr_t</code>	int	<code>omp_fr_cuda</code>
<code>fr_name</code>	<code>const char *</code>	str	<code>"cuda"</code>
<code>vendor</code>	int	int	11
<code>vendor_name</code>	<code>const char *</code>	str	<code>"nvidia"</code>
<code>device_num</code>	int	int	

platform	N/A	
device	int	int
device_context	N/A	
targetsync	cudaStream_t	ptr

Available properties for a CUDA driver API interop object:

Property	C data type	API routine	value (if constant)
fr_id	omp_interop_fr_t	int	omp_fr_cuda_driver
fr_name	const char *	str	"cuda_driver"
vendor	int	int	11
vendor_name	const char *	str	"nvidia"
device_num	int	int	
platform	N/A		
device	CUdevice	int	
device_context	CUcontext	ptr	
targetsync	CUstream	ptr	

Available properties for an HIP interop object:

Property	C data type	API routine	value (if constant)
fr_id	omp_interop_fr_t	int	omp_fr_hip
fr_name	const char *	str	"hip"
vendor	int	int	11
vendor_name	const char *	str	"nvidia"
device_num	int	int	
platform	N/A		
device	hipDevice_t	int	
device_context	hipCtx_t	ptr	
targetsync	hipStream_t	ptr	

13 The libgomp ABI

The following sections present notes on the external ABI as presented by libgomp. Only maintainers should need them.

13.1 Implementing MASKED and MASTER construct

```
if (omp_get_thread_num () == thread_num)
    block
```

Hereby, *thread_num* has the value of the argument to the `filter` clause or zero if not specified.

Alternately, we generate two copies of the parallel subfunction and only include this in the version run by the *thread_num* thread. Surely this is not worthwhile though...

13.2 Implementing CRITICAL construct

Without a specified name,

```
void GOMP_critical_start (void);
void GOMP_critical_end (void);
```

so that we don't get COPY relocations from libgomp to the main application.

With a specified name, use `omp_set_lock` and `omp_unset_lock` with name being transformed into a variable declared like

```
omp_lock_t gomp_critical_user_<name> __attribute__((common))
```

Ideally the ABI would specify that all zero is a valid unlocked state, and so we wouldn't need to initialize this at startup.

13.3 Implementing ATOMIC construct

The target should implement the `__sync` builtins.

Failing that we could add

```
void GOMP_atomic_enter (void)
void GOMP_atomic_exit (void)
```

which reuses the regular lock code, but with yet another lock object private to the library.

13.4 Implementing FLUSH construct

Expands to the `__sync_synchronize` builtin.

13.5 Implementing BARRIER construct

```
void GOMP_barrier (void)
```

13.6 Implementing THREADPRIVATE construct

In `_most_` cases we can map this directly to `__thread`. Except that OMP allows constructors for C++ objects. We can either refuse to support this (how often is it used?) or we can implement something akin to `.ctors`.

Even more ideally, this ctor feature is handled by extensions to the main pthreads library. Failing that, we can have a set of entry points to register ctor functions to be called.

13.7 Implementing PRIVATE clause

In association with a PARALLEL, or within the lexical extent of a PARALLEL block, the variable becomes a local variable in the parallel subfunction.

In association with FOR or SECTIONS blocks, create a new automatic variable within the current function. This preserves the semantic of new variable creation.

13.8 Implementing FIRSTPRIVATE LASTPRIVATE COPYIN and COPYPRIVATE clauses

This seems simple enough for PARALLEL blocks. Create a private struct for communicating between the parent and subfunction. In the parent, copy in values for scalar and "small" structs; copy in addresses for others TREE_ADDRESSABLE types. In the subfunction, copy the value into the local variable.

It is not clear what to do with bare FOR or SECTION blocks. The only thing I can figure is that we do something like:

```
#pragma omp for firstprivate(x) lastprivate(y)
for (int i = 0; i < n; ++i)
  body;
```

which becomes

```
{
  int x = x, y;

  // for stuff

  if (i == n)
    y = y;
}
```

where the "x=x" and "y=y" assignments actually have different uids for the two variables, i.e. not something you could write directly in C. Presumably this only makes sense if the "outer" x and y are global variables.

COPYPRIVATE would work the same way, except the structure broadcast would have to happen via SINGLE machinery instead.

13.9 Implementing REDUCTION clause

The private struct mentioned in the previous section should have a pointer to an array of the type of the variable, indexed by the thread's *team_id*. The thread stores its final value into the array, and after the barrier, the primary thread iterates over the array to collect the values.

13.10 Implementing PARALLEL construct

```
#pragma omp parallel
{
  body;
}
```

becomes

```
void subfunction (void *data)
{
```

```

    use data;
    body;
}

setup data;
GOMP_parallel_start (subfunction, &data, num_threads);
subfunction (&data);
GOMP_parallel_end ();

void GOMP_parallel_start (void (*fn)(void *), void *data, unsigned num_threads)

```

The *FN* argument is the subfunction to be run in parallel.

The *DATA* argument is a pointer to a structure used to communicate data in and out of the subfunction, as discussed above with respect to *FIRSTPRIVATE* et al.

The *NUM_THREADS* argument is 1 if an *IF* clause is present and false, or the value of the *NUM_THREADS* clause, if present, or 0.

The function needs to create the appropriate number of threads and/or launch them from the dock. It needs to create the team structure and assign team ids.

```
void GOMP_parallel_end (void)
```

Tears down the team and returns us to the previous `omp_in_parallel()` state.

13.11 Implementing FOR construct

```

#pragma omp parallel for
for (i = lb; i <= ub; i++)
    body;

```

becomes

```

void subfunction (void *data)
{
    long _s0, _e0;
    while (GOMP_loop_static_next (&_s0, &_e0))
    {
        long _e1 = _e0, i;
        for (i = _s0; i < _e1; i++)
            body;
    }
    GOMP_loop_end_nowait ();
}

GOMP_parallel_loop_static (subfunction, NULL, 0, lb, ub+1, 1, 0);
subfunction (NULL);
GOMP_parallel_end ();

#pragma omp for schedule(runtime)
for (i = 0; i < n; i++)
    body;

```

becomes

```

{
    long i, _s0, _e0;
    if (GOMP_loop_runtime_start (0, n, 1, &_s0, &_e0))
        do {
            long _e1 = _e0;
            for (i = _s0, i < _e0; i++)
                body;
        } while (GOMP_loop_runtime_next (&_s0, &_e0));
}

```

```
GOMP_loop_end ();
}
```

Note that while it looks like there is trickiness to propagating a non-constant STEP, there isn't really. We're explicitly allowed to evaluate it as many times as we want, and any variables involved should automatically be handled as PRIVATE or SHARED like any other variables. So the expression should remain evaluable in the subfunction. We can also pull it into a local variable if we like, but since its supposed to remain unchanged, we can also not if we like.

If we have SCHEDULE(STATIC), and no ORDERED, then we ought to be able to get away with no work-sharing context at all, since we can simply perform the arithmetic directly in each thread to divide up the iterations. Which would mean that we wouldn't need to call any of these routines.

There are separate routines for handling loops with an ORDERED clause. Bookkeeping for that is non-trivial...

13.12 Implementing ORDERED construct

```
void GOMP_ordered_start (void)
void GOMP_ordered_end (void)
```

13.13 Implementing SECTIONS construct

A block as

```
#pragma omp sections
{
    #pragma omp section
    stmt1;
    #pragma omp section
    stmt2;
    #pragma omp section
    stmt3;
}
```

becomes

```
for (i = GOMP_sections_start (3); i != 0; i = GOMP_sections_next ())
    switch (i)
    {
        case 1:
            stmt1;
            break;
        case 2:
            stmt2;
            break;
        case 3:
            stmt3;
            break;
    }
GOMP_barrier ();
```

13.14 Implementing SINGLE construct

A block like

```
#pragma omp single
```

```

{
    body;
}

```

becomes

```

if (GOMP_single_start ())
    body;
GOMP_barrier ();

```

while

```

#pragma omp single copyprivate(x)
    body;

```

becomes

```

datap = GOMP_single_copy_start ();
if (datap == NULL)
{
    body;
    data.x = x;
    GOMP_single_copy_end (&data);
}
else
    x = datap->x;
GOMP_barrier ();

```

13.15 Implementing OpenACC's PARALLEL construct

```

void GOACC_parallel ()

```


14 Reporting Bugs

Bugs in the GNU Offloading and Multi Processing Runtime Library should be reported via Bugzilla (<https://gcc.gnu.org/bugzilla/>). Please add "openacc", or "openmp", or both to the keywords field in the bug report, as appropriate.

GNU General Public License

Version 3, 29 June 2007

Copyright © 2007 Free Software Foundation, Inc. <https://www.fsf.org>

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The GNU General Public License is a free, copyleft license for software and other kinds of works.

The licenses for most software and other practical works are designed to take away your freedom to share and change the works. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change all versions of a program—to make sure it remains free software for all its users. We, the Free Software Foundation, use the GNU General Public License for most of our software; it applies also to any other work released this way by its authors. You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for them if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs, and that you know you can do these things.

To protect your rights, we need to prevent others from denying you these rights or asking you to surrender the rights. Therefore, you have certain responsibilities if you distribute copies of the software, or if you modify it: responsibilities to respect the freedom of others.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must pass on to the recipients the same freedoms that you received. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

Developers that use the GNU GPL protect your rights with two steps: (1) assert copyright on the software, and (2) offer you this License giving you legal permission to copy, distribute and/or modify it.

For the developers' and authors' protection, the GPL clearly explains that there is no warranty for this free software. For both users' and authors' sake, the GPL requires that modified versions be marked as changed, so that their problems will not be attributed erroneously to authors of previous versions.

Some devices are designed to deny users access to install or run modified versions of the software inside them, although the manufacturer can do so. This is fundamentally incompatible with the aim of protecting users' freedom to change the software. The systematic pattern of such abuse occurs in the area of products for individuals to use, which is precisely where it is most unacceptable. Therefore, we have designed this version of the GPL to prohibit the practice for those products. If such problems arise substantially in other domains, we stand ready to extend this provision to those domains in future versions of the GPL, as needed to protect the freedom of users.

Finally, every program is threatened constantly by software patents. States should not allow patents to restrict development and use of software on general-purpose computers, but in those that do, we wish to avoid the special danger that patents applied to a free program could make it effectively proprietary. To prevent this, the GPL assures that patents cannot be used to render the program non-free.

The precise terms and conditions for copying, distribution and modification follow.

TERMS AND CONDITIONS

0. Definitions.

“This License” refers to version 3 of the GNU General Public License.

“Copyright” also means copyright-like laws that apply to other kinds of works, such as semiconductor masks.

“The Program” refers to any copyrightable work licensed under this License. Each licensee is addressed as “you”. “Licensees” and “recipients” may be individuals or organizations.

To “modify” a work means to copy from or adapt all or part of the work in a fashion requiring copyright permission, other than the making of an exact copy. The resulting work is called a “modified version” of the earlier work or a work “based on” the earlier work.

A “covered work” means either the unmodified Program or a work based on the Program.

To “propagate” a work means to do anything with it that, without permission, would make you directly or secondarily liable for infringement under applicable copyright law, except executing it on a computer or modifying a private copy. Propagation includes copying, distribution (with or without modification), making available to the public, and in some countries other activities as well.

To “convey” a work means any kind of propagation that enables other parties to make or receive copies. Mere interaction with a user through a computer network, with no transfer of a copy, is not conveying.

An interactive user interface displays “Appropriate Legal Notices” to the extent that it includes a convenient and prominently visible feature that (1) displays an appropriate copyright notice, and (2) tells the user that there is no warranty for the work (except to the extent that warranties are provided), that licensees may convey the work under this License, and how to view a copy of this License. If the interface presents a list of user commands or options, such as a menu, a prominent item in the list meets this criterion.

1. Source Code.

The “source code” for a work means the preferred form of the work for making modifications to it. “Object code” means any non-source form of a work.

A “Standard Interface” means an interface that either is an official standard defined by a recognized standards body, or, in the case of interfaces specified for a particular programming language, one that is widely used among developers working in that language.

The “System Libraries” of an executable work include anything, other than the work as a whole, that (a) is included in the normal form of packaging a Major Component, but which is not part of that Major Component, and (b) serves only to enable use of the work with that Major Component, or to implement a Standard Interface for which an implementation is available to the public in source code form. A “Major Component”, in this context, means a major essential component (kernel, window system, and so on) of the specific operating system (if any) on which the executable work runs, or a compiler used to produce the work, or an object code interpreter used to run it.

The “Corresponding Source” for a work in object code form means all the source code needed to generate, install, and (for an executable work) run the object code and to modify the work, including scripts to control those activities. However, it does not include the work’s System Libraries, or general-purpose tools or generally available free programs which are used unmodified in performing those activities but which are not part of the work. For example, Corresponding Source includes interface definition files associated with source files for the work, and the source code for shared libraries and dynamically linked subprograms that the work is specifically designed to require, such as by intimate data communication or control flow between those subprograms and other parts of the work.

The Corresponding Source need not include anything that users can regenerate automatically from other parts of the Corresponding Source.

The Corresponding Source for a work in source code form is that same work.

2. Basic Permissions.

All rights granted under this License are granted for the term of copyright on the Program, and are irrevocable provided the stated conditions are met. This License explicitly affirms your unlimited permission to run the unmodified Program. The output from running a covered work is covered by this License only if the output, given its content, constitutes a covered work. This License acknowledges your rights of fair use or other equivalent, as provided by copyright law.

You may make, run and propagate covered works that you do not convey, without conditions so long as your license otherwise remains in force. You may convey covered works to others for the sole purpose of having them make modifications exclusively for you, or provide you with facilities for running those works, provided that you comply with the terms of this License in conveying all material for which you do not control copyright. Those thus making or running the covered works for you must do so exclusively on your behalf, under your direction and control, on terms that prohibit them from making any copies of your copyrighted material outside their relationship with you.

Conveying under any other circumstances is permitted solely under the conditions stated below. Sublicensing is not allowed; section 10 makes it unnecessary.

3. Protecting Users’ Legal Rights From Anti-Circumvention Law.

No covered work shall be deemed part of an effective technological measure under any applicable law fulfilling obligations under article 11 of the WIPO copyright treaty adopted on 20 December 1996, or similar laws prohibiting or restricting circumvention of such measures.

When you convey a covered work, you waive any legal power to forbid circumvention of technological measures to the extent such circumvention is effected by exercising rights under this License with respect to the covered work, and you disclaim any intention to limit operation or modification of the work as a means of enforcing, against the work's users, your or third parties' legal rights to forbid circumvention of technological measures.

4. Conveying Verbatim Copies.

You may convey verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice; keep intact all notices stating that this License and any non-permissive terms added in accord with section 7 apply to the code; keep intact all notices of the absence of any warranty; and give all recipients a copy of this License along with the Program.

You may charge any price or no price for each copy that you convey, and you may offer support or warranty protection for a fee.

5. Conveying Modified Source Versions.

You may convey a work based on the Program, or the modifications to produce it from the Program, in the form of source code under the terms of section 4, provided that you also meet all of these conditions:

- a. The work must carry prominent notices stating that you modified it, and giving a relevant date.
- b. The work must carry prominent notices stating that it is released under this License and any conditions added under section 7. This requirement modifies the requirement in section 4 to "keep intact all notices".
- c. You must license the entire work, as a whole, under this License to anyone who comes into possession of a copy. This License will therefore apply, along with any applicable section 7 additional terms, to the whole of the work, and all its parts, regardless of how they are packaged. This License gives no permission to license the work in any other way, but it does not invalidate such permission if you have separately received it.
- d. If the work has interactive user interfaces, each must display Appropriate Legal Notices; however, if the Program has interactive interfaces that do not display Appropriate Legal Notices, your work need not make them do so.

A compilation of a covered work with other separate and independent works, which are not by their nature extensions of the covered work, and which are not combined with it such as to form a larger program, in or on a volume of a storage or distribution medium, is called an "aggregate" if the compilation and its resulting copyright are not used to limit the access or legal rights of the compilation's users beyond what the individual works permit. Inclusion of a covered work in an aggregate does not cause this License to apply to the other parts of the aggregate.

6. Conveying Non-Source Forms.

You may convey a covered work in object code form under the terms of sections 4 and 5, provided that you also convey the machine-readable Corresponding Source under the terms of this License, in one of these ways:

- a. Convey the object code in, or embodied in, a physical product (including a physical distribution medium), accompanied by the Corresponding Source fixed on a durable physical medium customarily used for software interchange.
- b. Convey the object code in, or embodied in, a physical product (including a physical distribution medium), accompanied by a written offer, valid for at least three years and valid for as long as you offer spare parts or customer support for that product model, to give anyone who possesses the object code either (1) a copy of the Corresponding Source for all the software in the product that is covered by this License, on a durable physical medium customarily used for software interchange, for a price no more than your reasonable cost of physically performing this conveying of source, or (2) access to copy the Corresponding Source from a network server at no charge.
- c. Convey individual copies of the object code with a copy of the written offer to provide the Corresponding Source. This alternative is allowed only occasionally and noncommercially, and only if you received the object code with such an offer, in accord with subsection 6b.
- d. Convey the object code by offering access from a designated place (gratis or for a charge), and offer equivalent access to the Corresponding Source in the same way through the same place at no further charge. You need not require recipients to copy the Corresponding Source along with the object code. If the place to copy the object code is a network server, the Corresponding Source may be on a different server (operated by you or a third party) that supports equivalent copying facilities, provided you maintain clear directions next to the object code saying where to find the Corresponding Source. Regardless of what server hosts the Corresponding Source, you remain obligated to ensure that it is available for as long as needed to satisfy these requirements.
- e. Convey the object code using peer-to-peer transmission, provided you inform other peers where the object code and Corresponding Source of the work are being offered to the general public at no charge under subsection 6d.

A separable portion of the object code, whose source code is excluded from the Corresponding Source as a System Library, need not be included in conveying the object code work.

A “User Product” is either (1) a “consumer product”, which means any tangible personal property which is normally used for personal, family, or household purposes, or (2) anything designed or sold for incorporation into a dwelling. In determining whether a product is a consumer product, doubtful cases shall be resolved in favor of coverage. For a particular product received by a particular user, “normally used” refers to a typical or common use of that class of product, regardless of the status of the particular user or of the way in which the particular user actually uses, or expects or is expected to use, the product. A product is a consumer product regardless of whether the product has substantial commercial, industrial or non-consumer uses, unless such uses represent the only significant mode of use of the product.

“Installation Information” for a User Product means any methods, procedures, authorization keys, or other information required to install and execute modified versions of a covered work in that User Product from a modified version of its Corresponding Source.

The information must suffice to ensure that the continued functioning of the modified object code is in no case prevented or interfered with solely because modification has been made.

If you convey an object code work under this section in, or with, or specifically for use in, a User Product, and the conveying occurs as part of a transaction in which the right of possession and use of the User Product is transferred to the recipient in perpetuity or for a fixed term (regardless of how the transaction is characterized), the Corresponding Source conveyed under this section must be accompanied by the Installation Information. But this requirement does not apply if neither you nor any third party retains the ability to install modified object code on the User Product (for example, the work has been installed in ROM).

The requirement to provide Installation Information does not include a requirement to continue to provide support service, warranty, or updates for a work that has been modified or installed by the recipient, or for the User Product in which it has been modified or installed. Access to a network may be denied when the modification itself materially and adversely affects the operation of the network or violates the rules and protocols for communication across the network.

Corresponding Source conveyed, and Installation Information provided, in accord with this section must be in a format that is publicly documented (and with an implementation available to the public in source code form), and must require no special password or key for unpacking, reading or copying.

7. Additional Terms.

“Additional permissions” are terms that supplement the terms of this License by making exceptions from one or more of its conditions. Additional permissions that are applicable to the entire Program shall be treated as though they were included in this License, to the extent that they are valid under applicable law. If additional permissions apply only to part of the Program, that part may be used separately under those permissions, but the entire Program remains governed by this License without regard to the additional permissions.

When you convey a copy of a covered work, you may at your option remove any additional permissions from that copy, or from any part of it. (Additional permissions may be written to require their own removal in certain cases when you modify the work.) You may place additional permissions on material, added by you to a covered work, for which you have or can give appropriate copyright permission.

Notwithstanding any other provision of this License, for material you add to a covered work, you may (if authorized by the copyright holders of that material) supplement the terms of this License with terms:

- a. Disclaiming warranty or limiting liability differently from the terms of sections 15 and 16 of this License; or
- b. Requiring preservation of specified reasonable legal notices or author attributions in that material or in the Appropriate Legal Notices displayed by works containing it; or
- c. Prohibiting misrepresentation of the origin of that material, or requiring that modified versions of such material be marked in reasonable ways as different from the original version; or

- d. Limiting the use for publicity purposes of names of licensors or authors of the material; or
- e. Declining to grant rights under trademark law for use of some trade names, trademarks, or service marks; or
- f. Requiring indemnification of licensors and authors of that material by anyone who conveys the material (or modified versions of it) with contractual assumptions of liability to the recipient, for any liability that these contractual assumptions directly impose on those licensors and authors.

All other non-permissive additional terms are considered “further restrictions” within the meaning of section 10. If the Program as you received it, or any part of it, contains a notice stating that it is governed by this License along with a term that is a further restriction, you may remove that term. If a license document contains a further restriction but permits relicensing or conveying under this License, you may add to a covered work material governed by the terms of that license document, provided that the further restriction does not survive such relicensing or conveying.

If you add terms to a covered work in accord with this section, you must place, in the relevant source files, a statement of the additional terms that apply to those files, or a notice indicating where to find the applicable terms.

Additional terms, permissive or non-permissive, may be stated in the form of a separately written license, or stated as exceptions; the above requirements apply either way.

8. Termination.

You may not propagate or modify a covered work except as expressly provided under this License. Any attempt otherwise to propagate or modify it is void, and will automatically terminate your rights under this License (including any patent licenses granted under the third paragraph of section 11).

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.

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, you do not qualify to receive new licenses for the same material under section 10.

9. Acceptance Not Required for Having Copies.

You are not required to accept this License in order to receive or run a copy of the Program. Ancillary propagation of a covered work occurring solely as a consequence of using peer-to-peer transmission to receive a copy likewise does not require acceptance.

However, nothing other than this License grants you permission to propagate or modify any covered work. These actions infringe copyright if you do not accept this License. Therefore, by modifying or propagating a covered work, you indicate your acceptance of this License to do so.

10. Automatic Licensing of Downstream Recipients.

Each time you convey a covered work, the recipient automatically receives a license from the original licensors, to run, modify and propagate that work, subject to this License. You are not responsible for enforcing compliance by third parties with this License.

An “entity transaction” is a transaction transferring control of an organization, or substantially all assets of one, or subdividing an organization, or merging organizations. If propagation of a covered work results from an entity transaction, each party to that transaction who receives a copy of the work also receives whatever licenses to the work the party’s predecessor in interest had or could give under the previous paragraph, plus a right to possession of the Corresponding Source of the work from the predecessor in interest, if the predecessor has it or can get it with reasonable efforts.

You may not impose any further restrictions on the exercise of the rights granted or affirmed under this License. For example, you may not impose a license fee, royalty, or other charge for exercise of rights granted under this License, and you may not initiate litigation (including a cross-claim or counterclaim in a lawsuit) alleging that any patent claim is infringed by making, using, selling, offering for sale, or importing the Program or any portion of it.

11. Patents.

A “contributor” is a copyright holder who authorizes use under this License of the Program or a work on which the Program is based. The work thus licensed is called the contributor’s “contributor version”.

A contributor’s “essential patent claims” are all patent claims owned or controlled by the contributor, whether already acquired or hereafter acquired, that would be infringed by some manner, permitted by this License, of making, using, or selling its contributor version, but do not include claims that would be infringed only as a consequence of further modification of the contributor version. For purposes of this definition, “control” includes the right to grant patent sublicenses in a manner consistent with the requirements of this License.

Each contributor grants you a non-exclusive, worldwide, royalty-free patent license under the contributor’s essential patent claims, to make, use, sell, offer for sale, import and otherwise run, modify and propagate the contents of its contributor version.

In the following three paragraphs, a “patent license” is any express agreement or commitment, however denominated, not to enforce a patent (such as an express permission to practice a patent or covenant not to sue for patent infringement). To “grant” such a patent license to a party means to make such an agreement or commitment not to enforce a patent against the party.

If you convey a covered work, knowingly relying on a patent license, and the Corresponding Source of the work is not available for anyone to copy, free of charge and under the terms of this License, through a publicly available network server or other readily accessible means, then you must either (1) cause the Corresponding Source to be so

available, or (2) arrange to deprive yourself of the benefit of the patent license for this particular work, or (3) arrange, in a manner consistent with the requirements of this License, to extend the patent license to downstream recipients. “Knowingly relying” means you have actual knowledge that, but for the patent license, your conveying the covered work in a country, or your recipient’s use of the covered work in a country, would infringe one or more identifiable patents in that country that you have reason to believe are valid.

If, pursuant to or in connection with a single transaction or arrangement, you convey, or propagate by procuring conveyance of, a covered work, and grant a patent license to some of the parties receiving the covered work authorizing them to use, propagate, modify or convey a specific copy of the covered work, then the patent license you grant is automatically extended to all recipients of the covered work and works based on it.

A patent license is “discriminatory” if it does not include within the scope of its coverage, prohibits the exercise of, or is conditioned on the non-exercise of one or more of the rights that are specifically granted under this License. You may not convey a covered work if you are a party to an arrangement with a third party that is in the business of distributing software, under which you make payment to the third party based on the extent of your activity of conveying the work, and under which the third party grants, to any of the parties who would receive the covered work from you, a discriminatory patent license (a) in connection with copies of the covered work conveyed by you (or copies made from those copies), or (b) primarily for and in connection with specific products or compilations that contain the covered work, unless you entered into that arrangement, or that patent license was granted, prior to 28 March 2007.

Nothing in this License shall be construed as excluding or limiting any implied license or other defenses to infringement that may otherwise be available to you under applicable patent law.

12. No Surrender of Others’ Freedom.

If conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot convey a covered work so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not convey it at all. For example, if you agree to terms that obligate you to collect a royalty for further conveying from those to whom you convey the Program, the only way you could satisfy both those terms and this License would be to refrain entirely from conveying the Program.

13. Use with the GNU Affero General Public License.

Notwithstanding any other provision of this License, you have permission to link or combine any covered work with a work licensed under version 3 of the GNU Affero General Public License into a single combined work, and to convey the resulting work. The terms of this License will continue to apply to the part which is the covered work, but the special requirements of the GNU Affero General Public License, section 13, concerning interaction through a network will apply to the combination as such.

14. Revised Versions of this License.

The Free Software Foundation may publish revised and/or new versions of the GNU General Public 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.

Each version is given a distinguishing version number. If the Program specifies that a certain numbered version of the GNU General Public License “or any later version” applies to it, you have the option of following the terms and conditions either of that numbered version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of the GNU General Public License, you may choose any version ever published by the Free Software Foundation.

If the Program specifies that a proxy can decide which future versions of the GNU General Public License can be used, that proxy’s public statement of acceptance of a version permanently authorizes you to choose that version for the Program.

Later license versions may give you additional or different permissions. However, no additional obligations are imposed on any author or copyright holder as a result of your choosing to follow a later version.

15. Disclaimer of Warranty.

THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

16. Limitation of Liability.

IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MODIFIES AND/OR CONVEYS THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

17. Interpretation of Sections 15 and 16.

If the disclaimer of warranty and limitation of liability provided above cannot be given local legal effect according to their terms, reviewing courts shall apply local law that most closely approximates an absolute waiver of all civil liability in connection with the Program, unless a warranty or assumption of liability accompanies a copy of the Program in return for a fee.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively state the exclusion of warranty; and each file should have at least the “copyright” line and a pointer to where the full notice is found.

```
one line to give the program's name and a brief idea of what it does.
Copyright (C) year name of author
```

```
This program is free software: you can redistribute it and/or modify
it under the terms of the GNU General Public License as published by
the Free Software Foundation, either version 3 of the License, or (at
your option) any later version.
```

```
This program is distributed in the hope that it will be useful, but
WITHOUT ANY WARRANTY; without even the implied warranty of
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU
General Public License for more details.
```

```
You should have received a copy of the GNU General Public License
along with this program. If not, see https://www.gnu.org/licenses/.
```

Also add information on how to contact you by electronic and paper mail.

If the program does terminal interaction, make it output a short notice like this when it starts in an interactive mode:

```
program Copyright (C) year name of author
This program comes with ABSOLUTELY NO WARRANTY; for details type 'show w'.
This is free software, and you are welcome to redistribute it
under certain conditions; type 'show c' for details.
```

The hypothetical commands ‘show w’ and ‘show c’ should show the appropriate parts of the General Public License. Of course, your program’s commands might be different; for a GUI interface, you would use an “about box”.

You should also get your employer (if you work as a programmer) or school, if any, to sign a “copyright disclaimer” for the program, if necessary. For more information on this, and how to apply and follow the GNU GPL, see <https://www.gnu.org/licenses/>.

The GNU General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Lesser General Public License instead of this License. But first, please read <https://www.gnu.org/licenses/why-not-lgpl.html>.

GNU Free Documentation License

Version 1.3, 3 November 2008

Copyright © 2000, 2001, 2002, 2007, 2008 Free Software Foundation, Inc.

<https://www.fsf.org>

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

0. 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 non-commercially. 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.

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 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.

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 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

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.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 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 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.

5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 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.”

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.

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 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.

8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 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 4) to Preserve its Title (section 1) will typically require changing the actual title.

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.

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.

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 <https://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.

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 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.

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.

Funding Free Software

If you want to have more free software a few years from now, it makes sense for you to help encourage people to contribute funds for its development. The most effective approach known is to encourage commercial redistributors to donate.

Users of free software systems can boost the pace of development by encouraging for-a-fee distributors to donate part of their selling price to free software developers—the Free Software Foundation, and others.

The way to convince distributors to do this is to demand it and expect it from them. So when you compare distributors, judge them partly by how much they give to free software development. Show distributors they must compete to be the one who gives the most.

To make this approach work, you must insist on numbers that you can compare, such as, “We will donate ten dollars to the Frobnitz project for each disk sold.” Don’t be satisfied with a vague promise, such as “A portion of the profits are donated,” since it doesn’t give a basis for comparison.

Even a precise fraction “of the profits from this disk” is not very meaningful, since creative accounting and unrelated business decisions can greatly alter what fraction of the sales price counts as profit. If the price you pay is \$50, ten percent of the profit is probably less than a dollar; it might be a few cents, or nothing at all.

Some redistributors do development work themselves. This is useful too; but to keep everyone honest, you need to inquire how much they do, and what kind. Some kinds of development make much more long-term difference than others. For example, maintaining a separate version of a program contributes very little; maintaining the standard version of a program for the whole community contributes much. Easy new ports contribute little, since someone else would surely do them; difficult ports such as adding a new CPU to the GNU Compiler Collection contribute more; major new features or packages contribute the most.

By establishing the idea that supporting further development is “the proper thing to do” when distributing free software for a fee, we can assure a steady flow of resources into making more free software.

Copyright © 1994 Free Software Foundation, Inc.

Verbatim copying and redistribution of this section is permitted without royalty; alteration is not permitted.

Library Index

A

acc_get_property 74
 acc_get_property_string 74

E

Environment Variable ... 59, 60, 61, 62, 63, 64, 65,
 66, 67, 68

F

FDL, GNU Free Documentation License 139

I

Implementation specific setting.. 63, 65, 66, 68, 107