

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 |

| | | |
|--------|--|----|
| 3.1.16 | <code>omp_get_max_active_levels</code> – Current maximum number of active regions | 21 |
| 3.1.17 | <code>omp_get_level</code> – Obtain the current nesting level..... | 21 |
| 3.1.18 | <code>omp_get_ancestor_thread_num</code> – Ancestor thread ID ... | 21 |
| 3.1.19 | <code>omp_get_team_size</code> – Number of threads in a team..... | 22 |
| 3.1.20 | <code>omp_get_active_level</code> – Number of parallel regions..... | 22 |
| 3.2 | Thread Affinity Routines | 22 |
| 3.2.1 | <code>omp_get_proc_bind</code> – Whether threads may be moved between CPUs | 22 |
| 3.3 | Teams Region Routines..... | 23 |
| 3.3.1 | <code>omp_get_num_teams</code> – Number of teams | 23 |
| 3.3.2 | <code>omp_get_team_num</code> – Get team number | 23 |
| 3.3.3 | <code>omp_set_num_teams</code> – Set upper teams limit for teams construct..... | 23 |
| 3.3.4 | <code>omp_get_max_teams</code> – Maximum number of teams of teams region | 24 |
| 3.3.5 | <code>omp_set_teams_thread_limit</code> – Set upper thread limit for teams construct..... | 24 |
| 3.3.6 | <code>omp_get_thread_limit</code> – Maximum number of threads ... | 24 |
| 3.4 | Tasking Routines..... | 25 |
| 3.4.1 | <code>omp_get_max_task_priority</code> – Maximum priority value .. | 25 |
| 3.4.2 | <code>omp_in_explicit_task</code> – Whether a given task is an explicit task..... | 25 |
| 3.4.3 | <code>omp_in_final</code> – Whether in final or included task region .. | 25 |
| 3.5 | Resource Relinquishing Routines..... | 26 |
| 3.5.1 | <code>omp_pause_resource</code> – Release OpenMP resources on a device | 26 |
| 3.5.2 | <code>omp_pause_resource_all</code> – Release OpenMP resources on all devices | 26 |
| 3.6 | Device Information Routines..... | 27 |
| 3.6.1 | <code>omp_get_num_procs</code> – Number of processors online..... | 27 |
| 3.6.2 | <code>omp_set_default_device</code> – Set the default device for target regions | 27 |
| 3.6.3 | <code>omp_get_default_device</code> – Get the default device for target regions | 27 |
| 3.6.4 | <code>omp_get_num_devices</code> – Number of target devices | 28 |
| 3.6.5 | <code>omp_get_device_num</code> – Return device number of current device | 28 |
| 3.6.6 | <code>omp_get_device_from_uid</code> – Obtain the device number to a unique id..... | 28 |
| 3.6.7 | <code>omp_get_uid_from_device</code> – Obtain the unique id of a device..... | 29 |
| 3.6.8 | <code>omp_is_initial_device</code> – Whether executing on the host device..... | 29 |
| 3.6.9 | <code>omp_get_initial_device</code> – Return device number of initial device..... | 30 |
| 3.7 | Device Memory Routines..... | 30 |

| | | |
|--------|--|----|
| 3.7.1 | <code>omp_target_alloc</code> – Allocate device memory | 30 |
| 3.7.2 | <code>omp_target_free</code> – Free device memory | 31 |
| 3.7.3 | <code>omp_target_is_present</code> – Check whether storage is mapped .. | 31 |
| 3.7.4 | <code>omp_target_is_accessible</code> – Check whether memory is device accessible | 32 |
| 3.7.5 | <code>omp_target_memcpy</code> – Copy data between devices | 33 |
| 3.7.6 | <code>omp_target_memcpy_async</code> – Copy data between devices asynchronously | 34 |
| 3.7.7 | <code>omp_target_memcpy_rect</code> – Copy a subvolume of data between devices | 35 |
| 3.7.8 | <code>omp_target_memcpy_rect_async</code> – Copy a subvolume of data between devices asynchronously | 36 |
| 3.7.9 | <code>omp_target_memset</code> – Set bytes in device memory | 37 |
| 3.7.10 | <code>omp_target_memset</code> – Set bytes in device memory asynchronously | 38 |
| 3.7.11 | <code>omp_target_associate_ptr</code> – Associate a device pointer with a host pointer | 39 |
| 3.7.12 | <code>omp_target_disassociate_ptr</code> – Remove device–host pointer association | 40 |
| 3.7.13 | <code>omp_get_mapped_ptr</code> – Return device pointer to a host pointer | 40 |
| 3.8 | Lock Routines | 41 |
| 3.8.1 | <code>omp_init_lock</code> – Initialize simple lock | 41 |
| 3.8.2 | <code>omp_init_nest_lock</code> – Initialize nested lock | 41 |
| 3.8.3 | <code>omp_destroy_lock</code> – Destroy simple lock | 42 |
| 3.8.4 | <code>omp_destroy_nest_lock</code> – Destroy nested lock | 42 |
| 3.8.5 | <code>omp_set_lock</code> – Wait for and set simple lock | 42 |
| 3.8.6 | <code>omp_set_nest_lock</code> – Wait for and set nested lock | 43 |
| 3.8.7 | <code>omp_unset_lock</code> – Unset simple lock | 43 |
| 3.8.8 | <code>omp_unset_nest_lock</code> – Unset nested lock | 43 |
| 3.8.9 | <code>omp_test_lock</code> – Test and set simple lock if available | 44 |
| 3.8.10 | <code>omp_test_nest_lock</code> – Test and set nested lock if available .. | 44 |
| 3.9 | Timing Routines | 44 |
| 3.9.1 | <code>omp_get_wtick</code> – Get timer precision | 45 |
| 3.9.2 | <code>omp_get_wtime</code> – Elapsed wall clock time | 45 |
| 3.10 | Event Routine | 45 |
| 3.10.1 | <code>omp_fulfill_event</code> – Fulfill and destroy an OpenMP event .. | 45 |
| 3.11 | Interoperability Routines | 46 |
| 3.11.1 | <code>omp_get_num_interop_properties</code> – Get the number of implementation-specific properties | 46 |
| 3.11.2 | <code>omp_get_interop_int</code> – Obtain integer-valued interoperability property | 46 |
| 3.11.3 | <code>omp_get_interop_ptr</code> – Obtain pointer-valued interoperability property | 47 |
| 3.11.4 | <code>omp_get_interop_str</code> – Obtain string-valued interoperability property | 48 |

| | | |
|---------|--|----|
| 3.11.5 | <code>omp_get_interop_name</code> – Obtain the name of an interop_property value as string..... | 48 |
| 3.11.6 | <code>omp_get_interop_type_desc</code> – Obtain type and description to an interop_property | 49 |
| 3.11.7 | <code>omp_get_interop_rc_desc</code> – Obtain error string to an interop_rc error code..... | 49 |
| 3.12 | Memory Management Routines..... | 50 |
| 3.12.1 | <code>omp_init_allocator</code> – Create an allocator..... | 50 |
| 3.12.2 | <code>omp_destroy_allocator</code> – Destroy an allocator..... | 51 |
| 3.12.3 | <code>omp_set_default_allocator</code> – Set the default allocator.. | 51 |
| 3.12.4 | <code>omp_get_default_allocator</code> – Get the default allocator.. | 52 |
| 3.12.5 | <code>omp_alloc</code> – Memory allocation with an allocator..... | 52 |
| 3.12.6 | <code>omp_aligned_alloc</code> – Memory allocation with an allocator and alignment | 53 |
| 3.12.7 | <code>omp_free</code> – Freeing memory allocated with OpenMP routines..... | 54 |
| 3.12.8 | <code>omp_calloc</code> – Allocate nullified memory with an allocator.. | 54 |
| 3.12.9 | <code>omp_aligned_calloc</code> – Allocate aligned nullified memory with an allocator..... | 55 |
| 3.12.10 | <code>omp_realloc</code> – Reallocate memory allocated with OpenMP routines..... | 56 |
| 3.13 | Environment Display Routine | 57 |
| 3.13.1 | <code>omp_display_env</code> – print the initial ICV values..... | 57 |

4 OpenMP Environment Variables 59

| | | |
|------|---|----|
| 4.1 | <code>OMP_ALLOCATOR</code> – Set the default allocator | 59 |
| 4.2 | <code>OMP_AFFINITY_FORMAT</code> – Set the format string used for affinity display | 60 |
| 4.3 | <code>OMP_CANCELLATION</code> – Set whether cancellation is activated | 61 |
| 4.4 | <code>OMP_DISPLAY_AFFINITY</code> – Display thread affinity information .. | 61 |
| 4.5 | <code>OMP_DISPLAY_ENV</code> – Show OpenMP version and environment variables..... | 61 |
| 4.6 | <code>OMP_DEFAULT_DEVICE</code> – Set the device used in target regions ... | 61 |
| 4.7 | <code>OMP_DYNAMIC</code> – Dynamic adjustment of threads | 62 |
| 4.8 | <code>OMP_MAX_ACTIVE_LEVELS</code> – Set the maximum number of nested parallel regions | 62 |
| 4.9 | <code>OMP_MAX_TASK_PRIORITY</code> – Set the maximum priority..... | 62 |
| 4.10 | <code>OMP_NESTED</code> – Nested parallel regions..... | 63 |
| 4.11 | <code>OMP_NUM_TEAMS</code> – Specifies the number of teams to use by teams region..... | 63 |
| 4.12 | <code>OMP_NUM_THREADS</code> – Specifies the number of threads to use.... | 63 |
| 4.13 | <code>OMP_PROC_BIND</code> – Whether threads may be moved between CPUs.. | 64 |
| 4.14 | <code>OMP_PLACES</code> – Specifies on which CPUs the threads should be placed | 64 |
| 4.15 | <code>OMP_STACKSIZE</code> – Set default thread stack size | 65 |
| 4.16 | <code>OMP_SCHEDULE</code> – How threads are scheduled..... | 66 |

| | | |
|----------|---|-----------|
| 4.17 | OMP_TARGET_OFFLOAD – Controls offloading behavior | 66 |
| 4.18 | OMP_TEAMS_THREAD_LIMIT – Set the maximum number of threads imposed by teams | 66 |
| 4.19 | OMP_THREAD_LIMIT – Set the maximum number of threads | 67 |
| 4.20 | OMP_WAIT_POLICY – How waiting threads are handled | 67 |
| 4.21 | GOMP_CPU_AFFINITY – Bind threads to specific CPUs | 67 |
| 4.22 | GOMP_DEBUG – Enable debugging output | 68 |
| 4.23 | GOMP_STACKSIZE – Set default thread stack size | 68 |
| 4.24 | GOMP_SPINCOUNT – Set the busy-wait spin count | 68 |
| 4.25 | GOMP_RTEMS_THREAD_POOLS – Set the RTEMS specific thread pools | 69 |
| 5 | Enabling OpenACC | 71 |
| 6 | OpenACC Runtime Library Routines | 73 |
| 6.1 | acc_get_num_devices – Get number of devices for given device type | 73 |
| 6.2 | acc_set_device_type – Set type of device accelerator to use... 73 | |
| 6.3 | acc_get_device_type – Get type of device accelerator to be used | 73 |
| 6.4 | acc_set_device_num – Set device number to use | 74 |
| 6.5 | acc_get_device_num – Get device number to be used | 74 |
| 6.6 | acc_get_property – Get device property | 74 |
| 6.7 | acc_async_test – Test for completion of a specific asynchronous operation | 75 |
| 6.8 | acc_async_test_all – Tests for completion of all asynchronous operations | 76 |
| 6.9 | acc_wait – Wait for completion of a specific asynchronous operation | 76 |
| 6.10 | acc_wait_all – Waits for completion of all asynchronous operations | 76 |
| 6.11 | acc_wait_all_async – Wait for completion of all asynchronous operations | 77 |
| 6.12 | acc_wait_async – Wait for completion of asynchronous operations | 77 |
| 6.13 | acc_init – Initialize runtime for a specific device type | 77 |
| 6.14 | acc_shutdown – Shuts down the runtime for a specific device type | 78 |
| 6.15 | acc_on_device – Whether executing on a particular device ... 78 | |
| 6.16 | acc_malloc – Allocate device memory | 78 |
| 6.17 | acc_free – Free device memory | 79 |
| 6.18 | acc_copyin – Allocate device memory and copy host memory to it | 79 |
| 6.19 | acc_present_or_copyin – If the data is not present on the device, allocate device memory and copy from host memory | 80 |

| | | |
|----------|--|-----------|
| 6.20 | <code>acc_create</code> – Allocate device memory and map it to host memory. | 80 |
| 6.21 | <code>acc_present_or_create</code> – If the data is not present on the device, allocate device memory and map it to host memory. | 81 |
| 6.22 | <code>acc_copyout</code> – Copy device memory to host memory. | 82 |
| 6.23 | <code>acc_delete</code> – Free device memory. | 83 |
| 6.24 | <code>acc_update_device</code> – Update device memory from mapped host memory. | 84 |
| 6.25 | <code>acc_update_self</code> – Update host memory from mapped device memory. | 84 |
| 6.26 | <code>acc_map_data</code> – Map previously allocated device memory to host memory. | 85 |
| 6.27 | <code>acc_unmap_data</code> – Unmap device memory from host memory. ... | 85 |
| 6.28 | <code>acc_deviceptr</code> – Get device pointer associated with specific host address. | 86 |
| 6.29 | <code>acc_hostptr</code> – Get host pointer associated with specific device address. | 86 |
| 6.30 | <code>acc_is_present</code> – Indicate whether host variable / array is present on device. | 86 |
| 6.31 | <code>acc_memcpy_to_device</code> – Copy host memory to device memory. ... | 87 |
| 6.32 | <code>acc_memcpy_from_device</code> – Copy device memory to host memory. | 87 |
| 6.33 | <code>acc_memcpy_device</code> – Copy memory within a device. | 88 |
| 6.34 | <code>acc_attach</code> – Let device pointer point to device-pointer target. ... | 89 |
| 6.35 | <code>acc_detach</code> – Let device pointer point to host-pointer target. ... | 89 |
| 6.36 | <code>acc_get_current_cuda_device</code> – Get CUDA device handle. ... | 90 |
| 6.37 | <code>acc_get_current_cuda_context</code> – Get CUDA context handle. ... | 90 |
| 6.38 | <code>acc_get_cuda_stream</code> – Get CUDA stream handle. | 90 |
| 6.39 | <code>acc_set_cuda_stream</code> – Set CUDA stream handle. | 90 |
| 6.40 | <code>acc_prof_register</code> – Register callbacks. | 91 |
| 6.41 | <code>acc_prof_unregister</code> – Unregister callbacks. | 91 |
| 6.42 | <code>acc_prof_lookup</code> – Obtain inquiry functions. | 91 |
| 6.43 | <code>acc_register_library</code> – Library registration. | 91 |
| 7 | OpenACC Environment Variables | 93 |
| 7.1 | <code>ACC_DEVICE_TYPE</code> | 93 |
| 7.2 | <code>ACC_DEVICE_NUM</code> | 93 |
| 7.3 | <code>ACC_PROFLIB</code> | 93 |
| 8 | CUDA Streams Usage | 95 |
| 9 | OpenACC Library Interoperability | 97 |
| 9.1 | Introduction. | 97 |
| 9.2 | First invocation: NVIDIA CUBLAS library API | 97 |
| 9.3 | First invocation: OpenACC library API | 98 |
| 9.4 | OpenACC library and environment variables | 99 |

| | | |
|-----------|---|------------|
| 10 | OpenACC Profiling Interface | 101 |
| 10.1 | Implementation Status and Implementation-Defined Behavior | 101 |
| 11 | OpenMP-Implementation Specifics | 107 |
| 11.1 | Implementation-defined ICV Initialization | 107 |
| 11.2 | OpenMP Context Selectors | 107 |
| 11.3 | Memory allocation | 107 |
| 12 | Offload-Target Specifics | 113 |
| 12.1 | AMD Radeon (GCN) | 113 |
| 12.1.1 | OpenMP <code>interop</code> – Foreign-Runtime Support for AMD GPUs | 114 |
| 12.2 | nvptx | 115 |
| 12.2.1 | OpenMP <code>interop</code> – Foreign-Runtime Support for Nvidia GPUs | 117 |
| 13 | The libgomp ABI | 119 |
| 13.1 | Implementing MASKED and MASTER construct | 119 |
| 13.2 | Implementing CRITICAL construct | 119 |
| 13.3 | Implementing ATOMIC construct | 119 |
| 13.4 | Implementing FLUSH construct | 119 |
| 13.5 | Implementing BARRIER construct | 119 |
| 13.6 | Implementing THREADPRIVATE construct | 119 |
| 13.7 | Implementing PRIVATE clause | 120 |
| 13.8 | Implementing FIRSTPRIVATE LASTPRIVATE COPYIN and COPYPRIVATE clauses | 120 |
| 13.9 | Implementing REDUCTION clause | 120 |
| 13.10 | Implementing PARALLEL construct | 120 |
| 13.11 | Implementing FOR construct | 121 |
| 13.12 | Implementing ORDERED construct | 122 |
| 13.13 | Implementing SECTIONS construct | 122 |
| 13.14 | Implementing SINGLE construct | 122 |
| 13.15 | Implementing OpenACC's PARALLEL construct | 123 |
| 14 | Reporting Bugs | 125 |
| | GNU General Public License | 127 |
| | GNU Free Documentation License | 139 |
| | ADDENDUM: How to use this License for your documents | 146 |
| | Funding Free Software | 147 |
| | Library Index | 149 |

1 Enabling OpenMP

To activate the OpenMP extensions for C/C++ and Fortran, the compile-time flag `-fopenmp` must be specified. For C and C++, this enables the handling of the OpenMP directives using `#pragma omp` and the `[[omp::directive(...)]], [[omp::sequence(...)]]` and `[[omp::decl(...)]]` attributes. For Fortran, it enables for free source form the `!$omp` sentinel for directives and the `!$` conditional compilation sentinel and for fixed source form the `c$omp`, `*$omp` and `!$omp` sentinels for directives and the `c$`, `*$` and `!$` conditional compilation sentinels. The flag also arranges for automatic linking of the OpenMP runtime library (Chapter 3 [Runtime Library Routines], page 15).

The `-fopenmp-simd` flag can be used to enable a subset of OpenMP directives that do not require the linking of either the OpenMP runtime library or the POSIX threads library.

A complete description of all OpenMP directives may be found in the OpenMP Application Program Interface (<https://www.openmp.org>) manuals. See also Chapter 2 [OpenMP Implementation Status], page 3.

2 OpenMP Implementation Status

The `_OPENMP` preprocessor macro and Fortran's `openmp_version` parameter, provided by `omp_lib.h` and the `omp_lib` module, have the value 202111 (i.e. OpenMP 5.2).

2.1 OpenMP 4.5

The OpenMP 4.5 specification is fully supported.

2.2 OpenMP 5.0

New features listed in Appendix B of the OpenMP specification

| Description | Status | Comments |
|--|--------|--|
| Array shaping | N | |
| Array sections with non-unit strides in C and C++ | N | |
| Iterators | Y | |
| <code>metadirective</code> directive | Y | |
| <code>declare variant</code> directive | Y | |
| <code>target-offload-var</code> ICV and <code>OMP_TARGET_OFFLOAD</code> env variable | Y | |
| Nested-parallel changes to <code>max-active-levels-var</code> ICV | Y | |
| <code>requires</code> directive | Y | See also Chapter 12 [Offload-Target Specifics], page 113, |
| <code>teams</code> construct outside an enclosing target region | Y | |
| Non-rectangular loop nests | P | Full support for C/C++, partial for Fortran (PR110735 (https://gcc.gnu.org/PR110735)) |
| <code>!=</code> as relational-op in canonical loop form for C/C++ | Y | |
| <code>nonmonotonic</code> as default loop schedule modifier for worksharing-loop constructs | Y | |
| Collapse of associated loops that are imperfectly nested loops | Y | |
| Clauses <code>if</code> , <code>nontemporal</code> and <code>order(concurrent)</code> in <code>simd</code> construct | Y | |
| <code>atomic</code> constructs in <code>simd</code> | Y | |
| <code>loop</code> construct | Y | |
| <code>order(concurrent)</code> clause | Y | |
| <code>scan</code> directive and <code>in_scan</code> modifier for the <code>reduction</code> clause | Y | |
| <code>in_reduction</code> clause on <code>task</code> constructs | Y | |
| <code>in_reduction</code> clause on <code>target</code> constructs | P | <code>nowait</code> only stub |
| <code>task_reduction</code> clause with <code>taskgroup</code> | Y | |
| <code>task</code> modifier to <code>reduction</code> clause | Y | |

| | | |
|---|---|--|
| <code>affinity</code> clause to <code>task</code> construct | Y | Stub only |
| <code>detach</code> clause to <code>task</code> construct | Y | |
| <code>omp_fulfill_event</code> runtime routine | Y | |
| <code>reduction</code> and <code>in_reduction</code> clauses on <code>taskloop</code> and <code>taskloop simd</code> constructs | Y | |
| <code>taskloop</code> construct cancelable by <code>cancel</code> construct | Y | |
| <code>mutexinoutset</code> <i>dependence-type</i> for <code>depend</code> clause | Y | |
| Predefined memory spaces, memory allocators, allocator traits | Y | See also Section 11.3 [Memory allocation], page 107, |
| Memory management routines | Y | |
| <code>allocate</code> directive | P | C++ unsupported; see also Section 11.3 [Memory allocation], page 107, |
| <code>allocate</code> clause | P | Clause has no effect on <code>target</code> (PR113436 (https://gcc.gnu.org/PR113436)) |
| <code>use_device_addr</code> clause on <code>target</code> data | Y | |
| <code>ancestor</code> modifier on <code>device</code> clause | Y | |
| Implicit declare <code>target</code> directive | Y | |
| Discontiguous array section with <code>target update</code> construct | N | |
| C/C++'s lvalue expressions in <code>to</code> , <code>from</code> and <code>map</code> clauses | Y | |
| C/C++'s lvalue expressions in <code>depend</code> clauses | Y | |
| Nested <code>declare target</code> directive | Y | |
| Combined <code>master</code> constructs | Y | |
| <code>depend</code> clause on <code>taskwait</code> | Y | |
| Weak memory ordering clauses on <code>atomic</code> and <code>flush</code> construct | Y | |
| <code>hint</code> clause on the <code>atomic</code> construct | Y | Stub only |
| <code>depobj</code> construct and depend objects | Y | |
| Lock hints were renamed to synchronization hints | Y | |
| <code>conditional</code> modifier to <code>lastprivate</code> clause | Y | |
| Map-order clarifications | P | |
| <code>close map-type-modifier</code> | Y | |
| Mapping C/C++ pointer variables and to assign the address of device memory mapped by an array section | P | |
| Mapping of Fortran pointer and allocatable variables, including pointer and allocatable components of variables | Y | |
| <code>defaultmap</code> extensions | Y | |

| | | |
|---|---|-------------------------------------|
| <code>declare mapper</code> directive | P | Initial support and for C/C++, only |
| <code>omp_get_supported_active_levels</code> routine | Y | |
| Runtime routines and environment variables to display runtime thread affinity information | Y | |
| <code>omp_pause_resource</code> and <code>omp_pause_resource_all</code> runtime routines | Y | |
| <code>omp_get_device_num</code> runtime routine | Y | |
| OMPT interface | N | |
| OMPD interface | N | |

Other new OpenMP 5.0 features

| Description | Status | Comments |
|---------------------------------------|--------|----------|
| Supporting C++'s range-based for loop | Y | |

2.3 OpenMP 5.1

New features listed in Appendix B of the OpenMP specification

| Description | Status | Comments |
|---|--------|---------------------------|
| OpenMP directive as C++ attribute specifiers | Y | |
| <code>omp_all_memory</code> reserved locator | Y | |
| <i>target_device trait</i> in OpenMP Context | Y | |
| <code>target_device</code> selector set in context selectors | Y | |
| C/C++'s delimited <code>declare variant</code> directive: support elision of preprocessed code and interpret enclosed function definitions as variant functions | Y | |
| <code>declare variant</code> : new clauses <code>adjust_args</code> and <code>append_args</code> | Y | |
| <code>dispatch</code> construct | Y | |
| device-specific ICV settings with environment variables | Y | |
| <code>assume</code> and <code>assumes</code> directives | Y | |
| <code>nothing</code> directive | Y | |
| <code>error</code> directive | Y | |
| <code>masked</code> construct | Y | |
| <code>scope</code> directive | Y | |
| Loop transformation constructs | Y | |
| <code>strict</code> modifier in the <code>grainsize</code> and <code>num_tasks</code> clauses of the <code>taskloop</code> construct | Y | |
| <code>align</code> clause in <code>allocate</code> directive | P | Only C and Fortran |
| <code>align</code> modifier in <code>allocate</code> clause | Y | |
| <code>thread_limit</code> clause to <code>target</code> construct | Y | |
| <code>has_device_addr</code> clause to <code>target</code> construct | Y | |
| Iterators in <code>target update</code> motion clauses and <code>map</code> clauses | P | Limited support for C/C++ |
| Indirect calls to the device version of a procedure or function in <code>target</code> regions | Y | |

| | | |
|---|---|--|
| <code>interop</code> directive | Y | Cf. Chapter 12 [Offload-Target Specifics], page 113, |
| <code>omp_interop_t</code> object support in runtime routines | Y | |
| <code>nowait</code> clause in <code>taskwait</code> directive | Y | |
| Extensions to the <code>atomic</code> directive | Y | |
| <code>seq_cst</code> clause on a <code>flush</code> construct | Y | |
| <code>inoutset</code> argument to the <code>depend</code> clause | Y | |
| <code>private</code> and <code>firstprivate</code> argument to <code>default</code> clause in C and C++ | Y | |
| <code>present</code> argument to <code>defaultmap</code> clause | Y | |
| <code>omp_set_num_teams</code> , <code>omp_set_teams_thread_limit</code> , <code>omp_get_max_teams</code> , <code>omp_get_teams_thread_limit</code> runtime routines | Y | |
| <code>omp_target_is_accessible</code> runtime routine | Y | |
| <code>omp_target_memcpy_async</code> and <code>omp_target_memcpy_rect_async</code> runtime routines | Y | |
| <code>omp_get_mapped_ptr</code> runtime routine | Y | |
| <code>omp_calloc</code> , <code>omp_realloc</code> , <code>omp_aligned_alloc</code> and <code>omp_aligned_calloc</code> runtime routines | Y | |
| <code>omp_alloctrail_key_t</code> enum: <code>omp_atv_serialized</code> added, <code>omp_atv_default</code> changed | Y | |
| <code>omp_display_env</code> runtime routine | Y | |
| <code>ompt_scope_endpoint_t</code> enum: <code>ompt_scope_beginend</code> | N | |
| <code>ompt_sync_region_t</code> enum additions | N | |
| <code>ompt_state_t</code> enum: <code>ompt_state_wait_barrier_implementation</code> and <code>ompt_state_wait_barrier_teams</code> | N | |
| <code>ompt_callback_target_data_op_emi_t</code> , <code>ompt_callback_target_emi_t</code> , <code>ompt_callback_target_map_emi_t</code> and <code>ompt_callback_target_submit_emi_t</code> | N | |
| <code>ompt_callback_error_t</code> type | N | |
| <code>OMP_PLACES</code> syntax extensions | Y | |
| <code>OMP_NUM_TEAMS</code> and <code>OMP_TEAMS_THREAD_LIMIT</code> environment variables | Y | |

Other new OpenMP 5.1 features

| Description | Status | Comments |
|---|--------|----------|
| Support of strictly structured blocks in Fortran | Y | |
| Support of structured block sequences in C/C++ | Y | |
| <code>unconstrained</code> and <code>reproducible</code> modifiers on <code>order</code> clause | Y | |
| Support <code>begin/end declare target</code> syntax in C/C++ | Y | |

| | | |
|--|---|--|
| Pointer predetermined firstprivate getting initialized to address of matching mapped list item per 5.1, Sect. 2.21.7.2 | N | |
| For Fortran, diagnose placing declarative before/between <code>USE</code> , <code>IMPORT</code> , and <code>IMPLICIT</code> as invalid | N | |
| Optional comma between directive and clause in the <code>#pragma</code> form | Y | |
| <code>indirect</code> clause in <code>declare target</code> | Y | |
| <code>device_type(nohost)/device_type(host)</code> for variables | N | |
| <code>present</code> modifier to the <code>map</code> , <code>to</code> and <code>from</code> clauses | Y | |
| Changed interaction between <code>declare target</code> and OpenMP context | Y | |
| Dynamic selector support in <code>metadirective</code> | Y | |
| Dynamic selector support in <code>declare variant</code> | P | Fortran rejects non-constant expressions in dynamic selectors; C/C++ reject expressions using argument variables. (PR113904 (https://gcc.gnu.org/PR113904)) |

2.4 OpenMP 5.2

New features listed in Appendix B of the OpenMP specification

| Description | Status | Comments |
|---|--------|--|
| <code>omp_in_explicit_task</code> routine and <i>explicit-task-var</i> ICV | Y | |
| <code>omp/ompx/omx</code> sentinels and <code>omp_/ompx_</code> namespaces | N/A | warning for <code>ompx/omx</code> sentinels ¹ |
| Clauses on <code>end</code> directive can be on directive | Y | |
| <code>destroy</code> clause with <code>destroy-var</code> argument on <code>depobj</code> | Y | |
| Deprecation of no-argument <code>destroy</code> clause on <code>depobj</code> | N/A | undeprecated in OpenMP 6 |
| <code>linear</code> clause syntax changes and <code>step</code> modifier | Y | |
| Deprecation of minus operator for reductions | Y | |
| Deprecation of separating <code>map</code> modifiers without comma | Y | |

¹ The `ompx` sentinel as C/C++ pragma and C++ attributes are warned for with `-Wunknown-pragmas` (implied by `-Wall`) and `-Wattributes` (enabled by default), respectively; for Fortran free-source code, there is a warning enabled by default and, for fixed-source code, the `omx` sentinel is warned for with `-Wsurprising` (enabled by `-Wall`). Unknown clauses are always rejected with an error.

| | | |
|---|---|--|
| <code>declare mapper</code> with iterator and <code>present</code> modifiers | N | |
| If a matching mapped list item is not found in the data environment, the pointer retains its original value | Y | |
| New <code>enter</code> clause as alias for <code>to</code> on declare target directive | Y | |
| Deprecation of <code>to</code> clause on declare target directive | Y | |
| Extended list of directives permitted in Fortran pure procedures | Y | |
| New <code>allocators</code> directive for Fortran | Y | |
| Deprecation of <code>allocate</code> directive for Fortran allocatables/pointers | Y | |
| Optional paired <code>end</code> directive with <code>dispatch</code> | Y | |
| New <code>memspace</code> and <code>traits</code> modifiers for <code>uses_allocators</code> | P | Only predefined allocators |
| Deprecation of <code>traits</code> array following the <code>allocator_handle</code> expression in <code>uses_allocators</code> | Y | |
| New <code>otherwise</code> clause as alias for <code>default</code> on metadirectives | Y | |
| Deprecation of <code>default</code> clause on metadirectives | Y | Both <code>otherwise</code> and <code>default</code> are accepted without diagnostics. |
| Deprecation of delimited form of <code>declare target</code> | Y | |
| Reproducible semantics changed for <code>order(concurrent)</code> | N | |
| <code>allocate</code> and <code>firstprivate</code> clauses on <code>scope</code> | Y | |
| <code>ompt_callback_work</code> | N | |
| Default map-type for the <code>map</code> clause in <code>target enter/exit data</code> | Y | |
| New <code>doacross</code> clause as alias for <code>depend</code> with <code>source/sink</code> modifier | Y | |
| Deprecation of <code>depend</code> with <code>source/sink</code> modifier | Y | |
| <code>omp_cur_iteration</code> keyword | Y | |

Other new OpenMP 5.2 features

| Description | Status | Comments |
|--|--------|----------|
| For Fortran, optional comma between directive and clause | N | |
| Conforming device numbers and <code>omp_initial_device</code> and <code>omp_invalid_device</code> enum/PARAMETER | Y | |
| Initial value of <i>default-device-var</i> ICV with <code>OMP_TARGET_OFFLOAD=mandatory</code> | Y | |
| <code>all</code> as <i>implicit-behavior</i> for <code>defaultmap</code> | Y | |
| <i>interop-types</i> in any position of the modifier list for the <code>init</code> clause of the <code>interop</code> construct | Y | |

| | |
|--|---|
| Invoke virtual member functions of C++ objects created on the host device on other devices | N |
| <code>mapper</code> as map-type modifier in <code>declare mapper</code> | N |

2.5 OpenMP 6.0

New features listed in Appendix B of the OpenMP specification

| | | |
|--|-----|-------------------------------------|
| Features deprecated in versions 5.0, 5.1 and 5.2 were removed | N/A | Backward compatibility |
| Full support for C23 was added | P | |
| Full support for C++23 was added | P | |
| Full support for Fortran 2023 was added | P | |
| <code>_ALL</code> suffix to the device-scope environment variables | P | Host device number wrongly accepted |
| <code>num_threads</code> clause now accepts a list | N | |
| Abstract names added for <code>OMP_NUM_THREADS</code> , <code>OMP_THREAD_LIMIT</code> and <code>OMP_TEAMS_THREAD_LIMIT</code> | N | |
| Supporting increments with abstract names in <code>OMP_PLACES</code> | N | |
| Extension of <code>OMP_DEFAULT_DEVICE</code> and new <code>OMP_AVAILABLE_DEVICES</code> environment vars | N | |
| New <code>uid</code> trait for target devices and for <code>OMP_AVAILABLE_DEVICES</code> and <code>OMP_DEFAULT_DEVICE</code> | N | |
| New <code>OMP_THREADS_RESERVE</code> environment variable | N | |
| The <code>decl</code> attribute was added to the C++ attribute syntax | Y | |
| The OpenMP directive syntax was extended to include C23 attribute specifiers | Y | |
| Support for pure directives in Fortran's <code>do concurrent</code> | N | |
| All inarguable clauses take now an optional Boolean argument | N | |
| The <code>adjust_args</code> clause was extended to specify the argument by position and supports variadic arguments | N | |
| For Fortran, <i>locator list</i> can be also function reference with data pointer result | N | |
| Concept of <i>assumed-size arrays</i> in C and C++ | N | |
| <i>directive-name-modifier</i> accepted in all clauses | N | |
| Extension of <code>interop</code> operation of <code>append_args</code> , allowing all modifiers of the <code>init</code> clause | Y | |
| New argument-free version of <code>depobj</code> with repeatable clauses and the <code>init</code> clause | N | |
| Undeprecate omitting the argument to the <code>depend</code> clause of the argument version of the <code>depend</code> construct | Y | |
| For Fortran, atomic with <code>BLOCK</code> construct and, for C/C++, with unlimited curly braces supported | N | |

| | | | |
|--|---|--|---------------|
| For Fortran, atomic with pointer comparison | N | | |
| For Fortran, atomic with enum and enumeration types | N | | |
| For Fortran, atomic compare with storing the comparison result | N | | |
| Canonical loop sequences and new <code>looprange</code> clause | N | | |
| For Fortran, handling polymorphic types in data-sharing-attribute clauses | P | <code>private</code> | not supported |
| For Fortran, rejecting polymorphic types in data-mapping clauses | N | not diagnosed (and mostly unsupported) | |
| New <code>taskgraph</code> construct including <code>saved</code> modifier and <code>replayable</code> clause | N | | |
| <code>default</code> clause on the <code>target</code> directive and accepting variable categories | N | | |
| Semantic change regarding the reference count update with <code>use_device_ptr</code> and <code>use_device_addr</code> | N | | |
| Support for inductions | N | | |
| Reduction over private variables with <code>reduction</code> clause | N | | |
| Implicit reduction identifiers of C++ classes | N | | |
| New <code>init_complete</code> clause to the <code>scan</code> directive | N | | |
| <code>ref</code> modifier to the <code>map</code> clause | N | | |
| New <code>storage</code> map-type modifier; context-dependent <code>alloc</code> and <code>release</code> are aliases | N | | |
| Change of the <i>map-type</i> property from <i>ultimate</i> to <i>default</i> | N | | |
| <code>self</code> modifier to <code>map</code> and <code>self</code> as <code>defaultmap</code> argument | N | | |
| Mapping of <i>assumed-size arrays</i> in C, C++ and Fortran | N | | |
| <code>delete</code> as delete-modifier not as map type | N | | |
| For Fortran, the <code>automap</code> modifier to the <code>enter</code> clause of <code>declare_target</code> | N | | |
| <code>groupprivate</code> directive | N | | |
| <code>local</code> clause to <code>declare_target</code> directive | N | | |
| <code>part_size</code> allocator trait for <code>interleaved</code> allocator partitions | N | | |
| <code>pin_device</code> , <code>preferred_device</code> and <code>target_access</code> allocator traits | N | | |
| <code>access</code> allocator trait changes | N | | |
| New <code>partitioner</code> value to <code>partition</code> allocator trait | N | | |
| Semicolon-separated list to <code>uses_allocators</code> | N | | |
| New <code>need_device_addr</code> modifier to <code>adjust_args</code> clause | N | | |
| <code>interop</code> clause to <code>dispatch</code> | Y | | |
| Scope requirement changes for <code>declare_target</code> | N | | |
| <code>message</code> and <code>severity</code> clauses to <code>parallel</code> directive | N | | |

| | |
|--|---|
| <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 |

| | |
|--|---|
| omp_get_device_from_uid and omp_get_uid_from_device routines | Y |
| omp_get_device_num_teams, omp_set_device_num_teams, omp_get_device_teams_thread_limit, and omp_set_device_teams_thread_limit routines | N |
| omp_target_memset and omp_target_memset_async routines | Y |
| Fortran version of the interop runtime routines | Y |
| Routines for obtaining memory spaces/allocators for shared/device memory | N |
| omp_get_memspace_num_resources routine | N |
| omp_get_memspace_pagesize routine | N |
| omp_get_submemspace routine | N |
| omp_init_mempartitioner, omp_destroy_mempartitioner, omp_init_mempartition, omp_destroy_mempartition, omp_mempartition_set_part, omp_mempartition_get_user_data routines | N |
| Deprecation of the target_data_op, target, target_map and target_submit callbacks and as values that set_callback must return | N |
| ompt_target_data_transfer and ompt_target_data_transfer_async values in ompt_target_data_op_t enum | N |
| The values ompt_target_data_transfer_to_device, ompt_target_data_transfer_from_device, ompt_target_data_transfer_to_device_async and ompt_target_data_transfer_from_device_async of the target_data_op OMPT type were deprecated | N |
| ompt_get_buffer_limits OMPT routine | N |

Deprecated features, unless listed above

| | |
|--|---|
| Deprecation of omitting the optional white space to separate adjacent keywords in the directive-name in Fortran (fixed and free source form) | N |
| Deprecation of the combiner expression in the declare_reduction argument | N |
| Deprecation of the Fortran include file omp_lib.h | N |

Other new OpenMP 6.0 features

| | |
|---|---|
| Multi-word directives now use underscore by default | N |
| Relaxed Fortran restrictions to the aligned clause | N |
| Mapping lambda captures | N |
| New omp_pause_stop_tool constant for omp_pause_resource | N |

In Fortran (fixed and free source form), spaces between directive names are mandatory

Update of the map-type decay for mapping and declare_mapper

2.6 OpenMP Technical Report 14

Technical Report (TR) 14 is the first preview for OpenMP 6.1.

New features listed in Appendix B of the OpenMP specification

| | |
|---|---|
| The <code>depth</code> clause to <code>fuse</code> directive | N |
| The <code>attach</code> modifier to the <code>map</code> clause | N |
| The <code>dyn_groupprivate</code> clause and the <code>omp_get_dyn_groupprivate_ptr</code> , <code>omp_get_dyn_groupprivate_size</code> , and <code>omp_get_dyn_groupprivate_size</code> routines | N |
| <code>begin declare_variant</code> directive in Fortran | N |
| <code>grid</code> and <code>tile</code> modifier to the <code>size</code> clause | N |
| New <code>flatten</code> loop-transforming directive | N |
| <code>scaled</code> modifier to <code>simdlen</code> clause | N |
| New <code>omp_default_device</code> identifier as conforming device number | Y |
| Clarify when <code>omp_target_is_accessible</code> routine returns zero | N |

Deprecated features, unless listed above

| | |
|--|---|
| Deprecation of conditional-update-capture structured block without a capture statement | N |
|--|---|

3 OpenMP Runtime Library Routines

The runtime routines described here are defined by Section 18 of the OpenMP specification in version 5.2.

3.1 Thread Team Routines

Routines controlling threads in the current contention group. They have C linkage and do not throw exceptions.

3.1.1 `omp_set_num_threads` – Set upper team size limit

Description:

Specifies the number of threads used by default in subsequent parallel sections, if those do not specify a `num_threads` clause. The argument of `omp_set_num_threads` shall be a positive integer.

C/C++:

Prototype: `void omp_set_num_threads(int num_threads);`

Fortran:

Interface: `subroutine omp_set_num_threads(num_threads)
integer, intent(in) :: num_threads`

See also: Section 4.12 [OMP_NUM_THREADS], page 63, Section 3.1.2 [omp_get_num_threads], page 15, Section 3.1.3 [omp_get_max_threads], page 16,

Reference: OpenMP specification v4.5 (<https://www.openmp.org>), Section 3.2.1.

3.1.2 `omp_get_num_threads` – Size of the active team

Description:

Returns the number of threads in the current team. In a sequential section of the program `omp_get_num_threads` returns 1.

The default team size may be initialized at startup by the `OMP_NUM_THREADS` environment variable. At runtime, the size of the current team may be set either by the `NUM_THREADS` clause or by `omp_set_num_threads`. If none of the above were used to define a specific value and `OMP_DYNAMIC` is disabled, one thread per CPU online is used.

C/C++:

Prototype: `int omp_get_num_threads(void);`

Fortran:

Interface: `integer function omp_get_num_threads()`

See also: Section 3.1.3 [omp_get_max_threads], page 16, Section 3.1.1 [omp_set_num_threads], page 15, Section 4.12 [OMP_NUM_THREADS], page 63,

Reference: OpenMP specification v4.5 (<https://www.openmp.org>), Section 3.2.2.

3.1.3 `omp_get_max_threads` – Maximum number of threads of parallel region

Description:

Return the maximum number of threads used for the current parallel region that does not use the clause `num_threads`.

C/C++:

Prototype: `int omp_get_max_threads(void);`

Fortran:

Interface: `integer function omp_get_max_threads()`

See also: Section 3.1.1 [`omp_set_num_threads`], page 15, Section 3.1.6 [`omp_set_dynamic`], page 17, Section 3.3.6 [`omp_get_thread_limit`], page 24,

Reference: OpenMP specification v4.5 (<https://www.openmp.org>), Section 3.2.3.

3.1.4 `omp_get_thread_num` – Current thread ID

Description:

Returns a unique thread identification number within the current team. In a sequential parts of the program, `omp_get_thread_num` always returns 0. In parallel regions the return value varies from 0 to `omp_get_num_threads-1` inclusive. The return value of the primary thread of a team is always 0.

C/C++:

Prototype: `int omp_get_thread_num(void);`

Fortran:

Interface: `integer function omp_get_thread_num()`

See also: Section 3.1.2 [`omp_get_num_threads`], page 15, Section 3.1.18 [`omp_get_ancestor_thread_num`], page 21,

Reference: OpenMP specification v4.5 (<https://www.openmp.org>), Section 3.2.4.

3.1.5 `omp_in_parallel` – Whether a parallel region is active

Description:

This function returns `true` if currently running in parallel, `false` otherwise. Here, `true` and `false` represent their language-specific counterparts.

C/C++:

Prototype: `int omp_in_parallel(void);`

Fortran:

Interface: `logical function omp_in_parallel()`

Reference: OpenMP specification v4.5 (<https://www.openmp.org>), Section 3.2.6.

3.1.9 `omp_set_nested` – Enable/disable nested parallel regions

Description:

Enable or disable nested parallel regions, i.e., whether team members are allowed to create new teams. The function takes the language-specific equivalent of `true` and `false`, where `true` enables dynamic adjustment of team sizes and `false` disables it.

Enabling nested parallel regions also sets the maximum number of active nested regions to the maximum supported. Disabling nested parallel regions sets the maximum number of active nested regions to one.

Note that the `omp_set_nested` API routine was deprecated in the OpenMP specification 5.0 in favor of `omp_set_max_active_levels`.

C/C++:

Prototype: `void omp_set_nested(int nested);`

Fortran:

Interface: `subroutine omp_set_nested(nested)`
 `logical, intent(in) :: nested`

See also: Section 3.1.10 [`omp_get_nested`], page 18, Section 3.1.15 [`omp_set_max_active_levels`], page 20, Section 4.8 [`OMP_MAX_ACTIVE_LEVELS`], page 62, Section 4.10 [`OMP_NESTED`], page 63,

Reference: OpenMP specification v4.5 (<https://www.openmp.org>), Section 3.2.10.

3.1.10 `omp_get_nested` – Nested parallel regions

Description:

This function returns `true` if nested parallel regions are enabled, `false` otherwise. Here, `true` and `false` represent their language-specific counterparts.

The state of nested parallel regions at startup depends on several environment variables. If `OMP_MAX_ACTIVE_LEVELS` is defined and is set to greater than one, then nested parallel regions will be enabled. If not defined, then the value of the `OMP_NESTED` environment variable will be followed if defined. If neither are defined, then if either `OMP_NUM_THREADS` or `OMP_PROC_BIND` are defined with a list of more than one value, then nested parallel regions are enabled. If none of these are defined, then nested parallel regions are disabled by default.

Nested parallel regions can be enabled or disabled at runtime using `omp_set_nested`, or by setting the maximum number of nested regions with `omp_set_max_active_levels` to one to disable, or above one to enable.

Note that the `omp_get_nested` API routine was deprecated in the OpenMP specification 5.0 in favor of `omp_get_max_active_levels`.

C/C++:

Prototype: `int omp_get_nested(void);`

Fortran:

Interface: `logical function omp_get_nested()`

See also: Section 3.1.16 [omp-get-max-active-levels], page 21, Section 3.1.9 [omp-set-nested], page 18, Section 4.8 [OMP_MAX_ACTIVE_LEVELS], page 62, Section 4.10 [OMP_NESTED], page 63,

Reference: OpenMP specification v4.5 (<https://www.openmp.org>), Section 3.2.11.

3.1.11 omp_set_schedule – Set the runtime scheduling method

Description:

Sets the runtime scheduling method. The *kind* argument can have the value `omp_sched_static`, `omp_sched_dynamic`, `omp_sched_guided` or `omp_sched_auto`. Except for `omp_sched_auto`, the chunk size is set to the value of *chunk_size* if positive, or to the default value if zero or negative. For `omp_sched_auto` the *chunk_size* argument is ignored.

C/C++

Prototype: `void omp_set_schedule(omp_sched_t kind, int chunk_size);`

Fortran:

Interface: `subroutine omp_set_schedule(kind, chunk_size)`
 `integer(kind=omp_sched_kind) kind`
 `integer chunk_size`

See also: Section 3.1.12 [omp-get-schedule], page 19, Section 4.16 [OMP_SCHEDULE], page 66,

Reference: OpenMP specification v4.5 (<https://www.openmp.org>), Section 3.2.12.

3.1.12 omp_get_schedule – Obtain the runtime scheduling method

Description:

Obtain the runtime scheduling method. The *kind* argument is set to `omp_sched_static`, `omp_sched_dynamic`, `omp_sched_guided` or `omp_sched_auto`. The second argument, *chunk_size*, is set to the chunk size.

C/C++

Prototype: `void omp_get_schedule(omp_sched_t *kind, int *chunk_size);`

Fortran:

Interface: `subroutine omp_get_schedule(kind, chunk_size)`
 `integer(kind=omp_sched_kind) kind`
 `integer chunk_size`

See also: Section 3.1.11 [omp-set-schedule], page 19, Section 4.16 [OMP_SCHEDULE], page 66,

Reference: OpenMP specification v4.5 (<https://www.openmp.org>), Section 3.2.13.

3.1.16 `omp_get_max_active_levels` – Current maximum number of active regions

Description:

This function obtains the maximum allowed number of nested, active parallel regions.

C/C++

Prototype: `int omp_get_max_active_levels(void);`

Fortran:

Interface: `integer function omp_get_max_active_levels()`

See also: Section 3.1.15 [`omp_set_max_active_levels`], page 20, Section 3.1.20 [`omp_get_active_level`], page 22,

Reference: OpenMP specification v4.5 (<https://www.openmp.org>), Section 3.2.16.

3.1.17 `omp_get_level` – Obtain the current nesting level

Description:

This function returns the nesting level for the parallel blocks, which enclose the calling call.

C/C++

Prototype: `int omp_get_level(void);`

Fortran:

Interface: `integer function omp_level()`

See also: Section 3.1.20 [`omp_get_active_level`], page 22,

Reference: OpenMP specification v4.5 (<https://www.openmp.org>), Section 3.2.17.

3.1.18 `omp_get_ancestor_thread_num` – Ancestor thread ID

Description:

This function returns the thread identification number for the given nesting level of the current thread. For values of *level* outside zero to `omp_get_level` -1 is returned; if *level* is `omp_get_level` the result is identical to `omp_get_thread_num`.

C/C++

Prototype: `int omp_get_ancestor_thread_num(int level);`

Fortran:

Interface: `integer function omp_get_ancestor_thread_num(level)`
 `integer level`

See also: Section 3.1.17 [`omp_get_level`], page 21, Section 3.1.4 [`omp_get_thread_num`], page 16, Section 3.1.19 [`omp_get_team_size`], page 22,

Reference: OpenMP specification v4.5 (<https://www.openmp.org>), Section 3.2.18.

See also: Section 4.11 [OMP_NUM_TEAMS], page 63, Section 3.3.1 [omp_get_num_teams], page 23, Section 3.3.4 [omp_get_max_teams], page 24,

Reference: OpenMP specification v5.1 (<https://www.openmp.org>), Section 3.4.3.

3.3.4 omp_get_max_teams – Maximum number of teams of teams region

Description:

Return the maximum number of teams used for the teams region that does not use the clause `num_teams`.

C/C++:

Prototype: `int omp_get_max_teams(void);`

Fortran:

Interface: `integer function omp_get_max_teams()`

See also: Section 3.3.3 [omp_set_num_teams], page 23, Section 3.3.1 [omp_get_num_teams], page 23,

Reference: OpenMP specification v5.1 (<https://www.openmp.org>), Section 3.4.4.

3.3.5 omp_set_teams_thread_limit – Set upper thread limit for teams construct

Description:

Specifies the upper bound for number of threads that are available for each team created by the teams construct which does not specify a `thread_limit` clause. The argument of `omp_set_teams_thread_limit` shall be a positive integer.

C/C++:

Prototype: `void omp_set_teams_thread_limit(int thread_limit);`

Fortran:

Interface: `subroutine omp_set_teams_thread_limit(thread_limit)`
 `integer, intent(in) :: thread_limit`

See also: Section 4.18 [OMP_TEAMS_THREAD_LIMIT], page 66, Section 3.1.13 [omp_get_teams_thread_limit], page 20, Section 3.3.6 [omp_get_thread_limit], page 24,

Reference: OpenMP specification v5.1 (<https://www.openmp.org>), Section 3.4.5.

3.3.6 omp_get_thread_limit – Maximum number of threads

Description:

Return the maximum number of threads of the program.

C/C++:

Prototype: `int omp_get_thread_limit(void);`

Fortran:

Interface: logical function omp_in_final()

Reference: OpenMP specification v4.5 (<https://www.openmp.org>), Section 3.2.21.

3.5 Resource Relinquishing Routines

Routines releasing resources used by the OpenMP runtime. They have C linkage and do not throw exceptions.

3.5.1 omp_pause_resource – Release OpenMP resources on a device

Description:

Free resources used by the OpenMP program and the runtime library on and for the device specified by *device_num*; on success, zero is returned and non-zero otherwise.

The value of *device_num* must be a conforming device number. The routine may not be called from within any explicit region and all explicit threads that do not bind to the implicit parallel region have finalized execution.

C/C++:

Prototype: int omp_pause_resource(omp_pause_resource_t kind,
 int device_num);

Fortran:

Interface: integer function omp_pause_resource(kind,
 device_num)
 integer (kind=omp_pause_resource_kind) kind
 integer device_num

Reference: OpenMP specification v5.0 (<https://www.openmp.org>), Section 3.2.43.

3.5.2 omp_pause_resource_all – Release OpenMP resources on all devices

Description:

Free resources used by the OpenMP program and the runtime library on all devices, including the host. On success, zero is returned and non-zero otherwise.

The routine may not be called from within any explicit region and all explicit threads that do not bind to the implicit parallel region have finalized execution.

C/C++:

Prototype: int omp_pause_resource(omp_pause_resource_t kind);

Fortran:

Interface: integer function omp_pause_resource(kind)
 integer (kind=omp_pause_resource_kind) kind

See also: Section 3.5.1 [omp_pause_resource], page 26,

Reference: OpenMP specification v5.0 (<https://www.openmp.org>), Section 3.2.44.

C/C++

Prototype: `void *omp_target_alloc(size_t size, int device_num)`

Fortran:

Interface: `type(c_ptr) function omp_target_alloc(size,
device_num) bind(C)
use, intrinsic :: iso_c_binding, only: c_ptr, c_int,
c_size_t
integer(c_size_t), value :: size
integer(c_int), value :: device_num`

See also: Section 3.7.2 [omp_target_free], page 31, Section 3.7.11 [omp_target_associate_ptr],
page 39,

Reference: OpenMP specification v5.1 (<https://www.openmp.org>), Section 3.8.1

3.7.2 omp_target_free – Free device memory

Description:

This routine frees memory allocated by the `omp_target_alloc` routine. The `device_ptr` argument must be either a null pointer or a device pointer returned by `omp_target_alloc` for the specified `device_num`. The device number `device_num` must be a conforming device number.

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

C/C++

Prototype: `void omp_target_free(void *device_ptr, int
device_num)`

Fortran:

Interface: `subroutine omp_target_free(device_ptr, device_num)
bind(C)
use, intrinsic :: iso_c_binding, only: c_ptr, c_int
type(c_ptr), value :: device_ptr
integer(c_int), value :: device_num`

See also: Section 3.7.1 [omp_target_alloc], page 30, Section 3.7.12 [omp_target_disassociate_ptr],
page 40,

Reference: OpenMP specification v5.1 (<https://www.openmp.org>), Section 3.8.2

3.7.3 omp_target_is_present – Check whether storage is mapped

Description:

This routine tests whether storage, identified by the host pointer `ptr` is mapped to the device specified by `device_num`. If so, it returns a nonzero value and otherwise zero.

In GCC, this includes self mapping such that `omp_target_is_present` returns `true` when `device_num` specifies the host or when the host and the device share


```

        size, device_num) bind(C)
    use, intrinsic :: iso_c_binding, only: c_ptr,
    c_size_t, c_int
    type(c_ptr), value :: ptr
    integer(c_size_t), value :: size
    integer(c_int), value :: device_num

```

See also: Section 3.7.11 [omp_target_associate_ptr], page 39,

Reference: OpenMP specification v5.1 (<https://www.openmp.org>), Section 3.8.4

3.7.5 omp_target_memcpy – Copy data between devices

Description:

This routine copies *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*. The routine returns zero on success 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(void *dst,
    const void *src,
    size_t length,
    size_t dst_offset,
    size_t src_offset,
    int dst_device_num,
    int src_device_num)

```

Fortran:

Interface:

```

integer(c_int) function omp_target_memcpy( &
    dst, src, length, dst_offset, src_offset, &
    dst_device_num, src_device_num) 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

```

See also: Section 3.7.6 [omp_target_memcpy_async], page 34, Section 3.7.7 [omp_target_memcpy_rect], page 35,

Reference: OpenMP specification v5.1 (<https://www.openmp.org>), Section 3.8.5

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

3.7.7 `omp_target_memcpy_rect` – Copy a subvolume of data between devices

Description:

This routine copies a subvolume of data from the device identified by device number *src_device_num* to device *dst_device_num*. The array has *num_dims* dimensions and each array element has a size of *element_size* bytes. The *volume* array specifies how many elements per dimension are copied. The full sizes of the destination and source arrays are given by the *dst_dimensions* and *src_dimensions* arguments, respectively. The offset per dimension to the first element to be copied is given by the *dst_offset* and *src_offset* arguments. The routine returns zero on success and non-zero otherwise.

The OpenMP specification only requires that *num_dims* up to three is supported. In order to find implementation-specific maximally supported number of dimensions, the routine returns this value when invoked with a null pointer to both the *dst* and *src* arguments. As GCC supports arbitrary dimensions, it returns `INT_MAX`.

The device-number arguments must be conforming device numbers, the *src* and *dst* must be either both null pointers or all of the following must be fulfilled: *element_size* and *num_dims* must be positive and the *volume*, offset and dimension arrays must have at least *num_dims* dimensions.

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

C/C++

Prototype:

```
int omp_target_memcpy_rect(void *dst,
    const void *src,
    size_t element_size,
    int num_dims,
    const size_t *volume,
    const size_t *dst_offset,
    const size_t *src_offset,
    const size_t *dst_dimensions,
    const size_t *src_dimensions,
    int dst_device_num,
    int src_device_num)
```

Fortran:

Interface:

```
integer(c_int) function omp_target_memcpy_rect( &
    dst, src, element_size, num_dims, volume, &
    dst_offset, src_offset, dst_dimensions, &
    src_dimensions, dst_device_num, src_device_num)
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 :: element_size,
    dst_offset, src_offset
```

```
integer(c_size_t), value :: volume, dst_dimensions,
src_dimensions
integer(c_int), value :: num_dims, dst_device_num,
src_device_num
```

See also: Section 3.7.8 [omp_target_memcpy_rect_async], page 36, Section 3.7.5 [omp_target_memcpy], page 33, Chapter 12 [Offload-Target Specifics], page 113,

Reference: OpenMP specification v5.1 (<https://www.openmp.org>), Section 3.8.6

3.7.8 omp_target_memcpy_rect_async – Copy a subvolume of data between devices asynchronously

Description:

This routine copies asynchronously a subvolume of data from the device identified by device number *src_device_num* to device *dst_device_num*. The array has *num_dims* dimensions and each array element has a size of *element_size* bytes. The *volume* array specifies how many elements per dimension are copied. The full sizes of the destination and source arrays are given by the *dst_dimensions* and *src_dimensions* arguments, respectively. The offset per dimension to the first element to be copied is given by the *dst_offset* and *src_offset* arguments. 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 on success and non-zero otherwise.

The OpenMP specification only requires that *num_dims* up to three is supported. In order to find implementation-specific maximally supported number of dimensions, the routine returns this value when invoked with a null pointer to both the *dst* and *src* arguments. As GCC supports arbitrary dimensions, it returns INT_MAX.

The device-number arguments must be conforming device numbers, the *src* and *dst* must be either both null pointers or all of the following must be fulfilled: *element_size* and *num_dims* must be positive and the *volume*, offset and dimension arrays must have at least *num_dims* dimensions.

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

C/C++

Prototype:

```
int omp_target_memcpy_rect_async(void *dst,
const void *src,
size_t element_size,
int num_dims,
const size_t *volume,
const size_t *dst_offset,
const size_t *src_offset,
const size_t *dst_dimensions,
```

```

const size_t *src_dimensions,
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_rect_
  async( &
    dst, src, element_size, num_dims, volume, &
    dst_offset, src_offset, dst_dimensions, &
    src_dimensions, 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 :: element_size,
  dst_offset, src_offset
integer(c_size_t), value :: volume, dst_dimensions,
  src_dimensions
integer(c_int), value :: num_dims, dst_device_num,
  src_device_num
integer(c_int), value :: depobj_count
integer(omp_depend_kind), optional :: depobj_
  list(*)

```

See also: Section 3.7.7 [omp_target_memcpy_rect], page 35, Section 3.7.6 [omp_target_memcpy_async], page 34, Chapter 12 [Offload-Target Specifics], page 113,

Reference: OpenMP specification v5.1 (<https://www.openmp.org>), Section 3.8.8

3.7.9 omp_target_memset – Set bytes in device memory

Description:

This routine fills memory on the device identified by device number *device_num*. Starting from the device address *ptr*, the first *count* bytes are set to the value *val*, converted to `unsigned char`. If *count* is zero, the routine has no effect; if *ptr* is NULL, the behavior is unspecified. The function returns *ptr*.

The *device_num* must be a conforming device number and *ptr* must be a valid device pointer for that device. Running this routine in a `target` region except on the initial device is not supported.

C/C++

Prototype:

```

void *omp_target_memcpy(void *ptr,
  int val,
  size_t count,
  int device_num)

```

Fortran:

Interface:

```

type(c_ptr) function omp_target_memset( &
    ptr, val, count, device_num) bind(C)
use, intrinsic :: iso_c_binding, only: c_ptr,
    c_size_t, c_int
type(c_ptr), value :: ptr
integer(c_size_t), value :: count
integer(c_int), value :: val, device_num

```

See also: Section 3.7.10 [omp_target_memset_async], page 38,

Reference: OpenMP specification v6.0 (<https://www.openmp.org>), Section 25.8.1

3.7.10 omp_target_memset – Set bytes in device memory asynchronously

Description:

This routine fills memory on the device identified by device number *device_num*. Starting from the device address *ptr*, the first *count* bytes are set to the value *val*, converted to `unsigned char`. If *count* is zero, the routine has no effect; if *ptr* is `NULL`, the behavior is unspecified. 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 function returns *ptr*.

The *device_num* must be a conforming device number and *ptr* must be a valid device pointer for that device. Running this routine in a `target` region except on the initial device is not supported.

C/C++

Prototype:

```

void *omp_target_memcpy_async(void *ptr,
    int val,
    size_t count,
    int device_num,
    int depobj_count,
    omp_depend_t *depobj_list)

```

Fortran:

Interface:

```

type(c_ptr) function omp_target_memset_async( &
    ptr, val, count, 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 :: ptr
integer(c_size_t), value :: count
integer(c_int), value :: val, device_num, depobj_
count
integer(omp_depend_kind), optional :: depobj_
list(*)

```

See also: Section 3.7.9 [omp_target_memset], page 37,

Reference: OpenMP specification v6.0 (<https://www.openmp.org>), Section 25.8.2

3.7.11 omp_target_associate_ptr – Associate a device pointer with a host pointer

Description:

This routine associates storage on the host with storage on a device identified by *device_num*. The device pointer is usually obtained by calling `omp_target_alloc` or by other means (but not by using the `map` clauses or the `declare target` directive). The host pointer should point to memory that has a storage size of at least *size*.

The *device_offset* parameter specifies the offset into *device_ptr* that is used as the base address for the device side of the mapping; the storage size should be at least *device_offset* plus *size*.

After the association, the host pointer can be used in a `map` clause and in the `to` and `from` clauses of the `target update` directive to transfer data between the associated pointers. The reference count of such associated storage is infinite. The association can be removed by calling `omp_target_disassociate_ptr` which should be done before the lifetime of either storage ends.

The routine returns nonzero (EINVAL) when the *device_num* is invalid, for when the initial device or the associated device shares memory with the host. `omp_target_associate_ptr` returns zero if *host_ptr* points into already associated storage that is fully inside of a previously associated memory. Otherwise, if the association was successful zero is returned; if none of the cases above apply, nonzero (EINVAL) is returned.

The `omp_target_is_present` routine can be used to test whether associated storage for a device pointer exists.

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

C/C++

Prototype:

```
int omp_target_associate_ptr(const void *host_ptr,
                             const void *device_ptr,
                             size_t size,
                             size_t device_offset,
                             int device_num)
```

Fortran:

Interface:

```
integer(c_int) function omp_target_associate_ptr(host_ptr, &
  device_ptr, size, device_offset, device_num)
  bind(C)
  use, intrinsic :: iso_c_binding, only: c_ptr, c_int,
  c_size_t
  type(c_ptr), value :: host_ptr, device_ptr
```


4.25 GOMP_RTEMS_THREAD_POOLS – Set the RTEMS specific thread pools

Description:

This environment variable is only used on the RTEMS real-time operating system. It determines the scheduler instance specific thread pools. The format for GOMP_RTEMS_THREAD_POOLS is a list of optional `<thread-pool-count>[$<priority>]@<scheduler-name>` configurations separated by `:` where:

- `<thread-pool-count>` is the thread pool count for this scheduler instance.
- `$<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 `IO`, `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 `IO`. 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.

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.

6.40 `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.

- 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`

- 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> |

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* the AMD GPUs present have the `HSA_AMD_SYSTEM_INFO_SVM_ACCESSIBLE_BY_DEFAULT` property; some systems require the "xnack" feature enabled for this to be true, in which case the runtime will attempt to set the `HSA_XNACK` environment variable to '1' automatically (user-set values are not overridden, and the setting only affects the executable itself and any child processes). If any AMD GPU device is not supported, all AMD GPUs 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 on the host with the `ompx_gnu_managed_mem_alloc` allocator or in the `ompx_gnu_managed_mem_space` (both GNU extensions) allocate memory equivalent to HIP Managed Memory, although *not* actually allocated using `hipMallocManaged`. 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. The ROCm runtime will automatically handle data migration between host and device as needed. Not all AMD GPU devices support this feature, and many that do require that `-mxnack=on` is configured at compile time. 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. If the current device does not support Unified Shared Memory (or it is not enabled with `HSA_XNACK=1`) then Managed Memory might still work, but allocations may only be visible to a single device (whichever was the default device when the *first* allocation was made).
- 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

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>.

- 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

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.

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, 69

F

FDL, GNU Free Documentation License 139

I

Implementation specific setting .. 63, 66, 68, 69, 107