



OpenMP, OpenACC & Offloading BoF

Tobias Burnus & Jakub Jelinek

GNU Tools @ LPC 2020

24–28 August 2020



Topics

■ OpenMP

- OpenMP 4.5 (2015-11,368p); supported since GCC 6, Fortran only partially
- OpenMP 5.0 (2018-11,666p); OpenMP 5.1 upcoming: TR8 ('19), TR9 ('20)
- Partially supported since GCC 9

■ OpenACC

- OpenACC 2.6 (2017-11,129p), supported since GCC 10;
OpenACC 2.7 ('18), 3.0 ('19), 3.1 ('20?)

Topics

■ Offloading

- NVPTX (since GCC 6), AMD GCN (since GCC 9/10)
- Xeon Phi (since GCC 5), ~~HSA/L (GCC 6 to 10)~~

■ Related Topics

- Other offloading, other concurrency (C++'s, Fortran's, Ada, ...)

OpenMP – Missing features

- Mapping changes (C++ this, Fortran allocatable components, array shaping), declare mapper
- Allocator clause/directive, unified shared address/memory
- Allocators high bandwidth memory support using memkind
- Ancestor modifier on device clause, reverse offloading
- Support for non-perfectly nested loops
- Metadirectives
- Declare variant (partially supported)
- OMPT + OMPD

OpenMP – Missing features

- Detach clause, `omp_fulfill_event`
- Lvalue expressions in `map/to/from` clauses
- The `omp_get_device_num` API
- `OMP_TARGET_OFFLOAD` env var
- Figure out NUMA topology and handle host teams accordingly

Offloading Performance

- OpenACC: “kernels”
 - Autoparallelization issues
- Optimizations:
 - Value propagation into offloading function
 - Avoid copying-out of values if not modified
 - Avoid copying-in if not changed
 - Early optimizations → avoid entries in offload table
vs. late (IPA, LTO) → cannot modify offload table (not using, pass NULL?)
 - Alias analysis

Offloading Targets

■ General

- Support more than one ISA; esp. GCN has incompatible `-march=`
Other software distribution changes?

■ nvptx

- Currently only `sm_30` + `sm_35` (current: `sm_80` of CUDA 11)

■ GCN

- Currently, `fiji`, `gfx900`, `gfx906`
- Debugging support

