

[Enter Post Title Here]

I attended Supercomputing in my third year as OpenMP CEO to both represent IBM and OpenMP. This was a big year for us as we closed with many milestones in what I call a Significant Paradigm shift in Parallelism. The most significant milestone was that the OpenMP Consortium has released OpenMP 4.0 in 2013 with new parallelism features that are productive, portable, and performant across C, C++, and Fortran. OpenMP 4.0 contains significant additions for accelerators, standardized for a broad set of architectures, and an industry-first support for SIMD vectorization. It was being showcased at SC13. In my previous [post](#), I gave an update on the state of OpenMP 4.0. Now it is ratified and implementation is well on its way.

The OpenMP ARB Consortium now has 26 members and is still growing, adding three new members in the last year:

- Red Hat/GCC
- Barcelona SuperComputing Centre
- University of Houston

Coming implementations of OpenMP 4.0 include GNU, and the Intel 13.1 compiler with support for accelerators. Clang has started with support with OpenMP 3.1.

Another major shift coming will be our mission statement which the members have been working on for a year now. It supports a move towards parallelism for a broader field, heterogeneous computing, and multiple memory models, while retaining strong support for shared memory and HPC.

At SC13, I gave the opening keynote at the OpenMP Birds of a Feather (BoF) session. This videos are [online](#). In this session, I introduced the upcoming International Workshop on OpenMP (IWOMP) for 2014, which will be held in Salvador, Brazil, hosted by Cray and a local group that specializes in computer education in Brazil. The location promises to be wonderful and there is significant OpenMP interest in Brazil as well as South America in general. Although early, the conference will generally be in September coinciding with Spring in the southern Hemisphere. The cutoff for submissions will be in May 2014, so please be prepared. In general, we have moved IWOMP from June previously to September starting with this year in Canberra Australia.

Our OpenMP staff ran the OpenMP show booth with many visitors inquiring about OpenMP which is used to support the current fastest [Top-500](#) Supercomputer, the Tianhe 2.

The BoF showcased the 4.0 Specification by the Chair of the OpenMP Language committee LLNL's CTO, Bronis de Supinski, as well as a humorous look at OpenWound by Oracle's Ruud van De Pas. This was followed with a look at the Clang implementation of OpenMP and its status which is supported by an Intel open sourced runtime, by Intel's Jim Cownie. This motivated a flurry of discussion about implementation status of compilers for OpenMP 4.0 and currently GNU and Intel already have significant implementation ready for release by mid 2014. Finally, we had a talk and demonstration of the HPC ToolsKit by Professor John Mellor-Crummey of Rice University on the upcoming Technical Report on A common Tools API for Profilers and Debuggers.

I also gave a 30 minute Exhibitor's Forum talk on OpenMP 4.0, and the future of OpenMP to about 40 people. The main content of my message is that OpenMP is now more agile than ever. We did that with the introduction of a Technical Report process that enables us to publish work in progress and give users a look ahead and implementer a chance to verify feasibility. This is actually a fairly common facility in most Standards. We continue to do that now with more frequent face-to-face meetings, that is longer and enables more work to be done. Next, I showed what kind of future features we are preparing for the next release of OpenMP:

- OpenMP Tools: Profilers and Debuggers
- Consumer style parallelism: event/async/futures
- Enhance Accelerator support
- Additional Looping constructs
- Transactional Memory, Speculative Execution
- Task Model refinements
- CPU Affinity
- Common Array Shaping
- Full Error Model
- Interoperability

- Rebase to new C/C++/Fortran Standards

This is clearly a very broad set of parallelism capabilities than what we have been traditionally been involved in. Well, we are moving forward too beyond our traditional involvement with shared memory parallelism. We have already started addressing non-shared memory architecture with accelerator support in 4.0. Now we are looking to reach beyond traditional High Performance Computing.

OpenMP is a living language, and as such will continue to grow. OpenMP is more agile, merging additional features such as a common tool support, more affinity, error model, interoperability with other models, new forms of loop parallelism, additional support for tasks, accelerators, event-driven programming, transactional memory, speculative execution and rebasing for new base language standards. The OpenMP revision cycle is increasing in speed and in predictability, while delivering concurrent technical reports and language extensions.