



News

[Press Releases](#)
[Event Announcements](#)
[CCT Weekly](#)
[Grants and Funding](#)
[Student News](#)
[Archived News](#)

LSU CCT's STE||AR Group to announce the sixth formal release of HPX (V0.9.6)

Louisiana State University's Center for Computation & Technology and its STE||AR Group are proud to announce the sixth formal release of HPX (V0.9.6).

High Performance ParallelX (HPX) provides a unified programming model for parallel and distributed applications of any scale. It is a general purpose C++ runtime system targeted at conventional, widely available architectures. In addition, HPX is the first freely available, open source, feature-complete, modular, and performance oriented implementation of the ParallelX execution model.

With the changes below, HPX is leading the charge of a whole new era of computation. By intrinsically breaking down and synchronizing the work to be done, HPX insures that application developers will no longer have to fret about where a segment of code executes. HPX allows coders to focus their time and energy to understanding the data dependencies of their algorithms and thereby the core obstacles to an efficient code.

Here are some of the advantages of using HPX:

- HPX exposes an API equivalent to the facilities as standardized by C++11/14 extended to distributed computing. Everything programmers know about primitives in the standard C++ library is still valid in the context of HPX.
- There is no need for the programmer to worry about lower level parallelization paradigms like threads or message passing; no need to understand pthreads, MPI, OpenMP, or Windows threads, etc.
- There is no need to think about different types of parallelism such as tasks, pipelines, or fork-join, task or data parallelism.
- The same source of your program compiles and runs on Linux, MacOS, Windows, and Android.
- The same code runs on shared memory multi-core systems and supercomputers, on handheld devices and Xeon-Phi accelerators, or a heterogeneous mix of those.

In this release we have made several significant changes:

- Consolidated API to be aligned with the C++11 Standard
- Implemented a distributed version of our Active Global Address Space (AGAS)
- Ported HPX to the Xeon-Phi device
- Added support for the SLURM scheduling system
- Improved the performance counter framework
- Added a distributed version of the Active Global Address Space (AGAS)
- Added parcel (message) compression and parcel coalescing systems
- Allow different scheduling policies for different parts of code with experimental executors API
- Added experimental security support on the locality level
- Created a native transport layer on top of Infiniband networks
- Created a native transport layer on top of low level MPI functions
- Added an experimental tuple-space object

We hope you will try out V0.9.6 and begin to contemplate where HPX can take your applications to the next level.

You can download the release from our website at <http://stellar.cct.lsu.edu/downloads/>, or get HPX directly from GitHub at <https://github.com/STELLAR-GROUP/hpx>. If you have suggestions, questions, or ideas we would love to hear from you. You can find us at our website, reach us at hpx-users@stellar.cct.lsu.edu, or chat with us live on IRC in the #stellar chat room on Freenode.

Publish Date:
08-02-2013

