



## LSU Delivers V0.9.7 of the HPX Runtime System

### News

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

The LSU Center for Computation & Technology's STE||AR Group is proud to announce the next formal release of HPX (V0.9.7). HPX is a general purpose C++ runtime system for parallel and distributed applications of any scale. This release focusses on performance, portability improvements, and better conformance to the C++11 and the upcoming C++14 standards.

In today's world of HPC, application developers struggle with the dilemma of spending hours finely tuning a complex, architecture specific code or writing a simple yet inefficient program which fails to take advantage the computing resources available. The STE||AR Group, has focused on bringing a reasonable solution to this unnecessary conundrum in the latest HPX release.

HPX is a novel combination of well-known ideas with new unique overarching concepts. It aims to resolve the problems related to scalability, resiliency, power efficiency, and runtime adaptive resource management that will be of growing importance as HPC architectures evolve from peta- to exa-scale. It departs from today's prevalent programming models with the goal of mitigating their respective limitations, such as implicit and explicit global barriers, coarse grain parallelism, and lack of overlap between computation and communication. It exposes a coherent programming model, unifying all different types of parallel execution in HPC systems. While it is not based on any conventional programming model such as MPI, OpenMP, CUDA, or OpenCL, it does integrate with those and provides a smooth migration path towards future computer architectures.

HPX represents an innovative mixture of global system-wide address space, fine grain parallelism, lightweight synchronization and work-queue based implicit message driven computation, full semantic equivalence of local and remote execution, and explicit support for hardware accelerators.

The focus of this release is on the portability of the runtime system: portability with the C++ standard and with a wide range of computer architectures. The API of HPX is C++11 compliant and is kept in line with the upcoming C++14 standard. Programmers who are familiar with C++11 will be able to effortlessly write code using HPX. HPX has been ported to new platforms. It is now supported on Linux, Mac OSX, Blue Gene/Q, and Windows. This wide range of supported platforms allows HPX to run on almost all of the top 500 machines, which translates to almost all HPX users having access to a supercomputer that they can run their code on. Additionally, HPX supports the Intel Xeon Phi® accelerator. The same code that a scientist writes to run on the cores of a machine can now be compiled for Intel's Xeon Phi® which gives the programmer access to the computing power of this exciting new accelerator with no extra effort.

"The HPX V0.9.7 release provides the basis of a solution which will allow users to write simple, efficient, and portable code. These ideas will change the way applications are written and will one day be pervasive throughout the HPC community" said Hartmut Kaiser, lead of the STE||AR group and adjunct associate research professor of LSU's Department of Computer Science. "HPX's portability and the compliance of the API to C++ lays the groundwork for its wide usability in different application domains."

For more information about HPX, as well as downloads and release notes, please visit <http://stellar.cct.lsu.edu>.

**Publish Date:**  
11-14-2013

