



# HPX Applications and Performance Adaptation

Alice Koniges<sup>1</sup>, Jayashree Ajay Candada<sup>2</sup>, Hartmut Kaiser<sup>3</sup>, Kevin Huck<sup>4</sup>, Jeremy Kemp<sup>5</sup>, Thomas Heller<sup>6</sup>, Matthew Anderson<sup>2</sup>, Andrew Lumsdaine<sup>2</sup>, Adrian Serio<sup>3</sup>, Michael Wolf<sup>7</sup>, Bryce Adelstein Lelbach<sup>1</sup>, Ron Brightwell<sup>7</sup>, Thomas Sterling<sup>2</sup>



The HPX runtime system is a critical component of the DOE XPRESS (eXascale PRogramming Environment and System Software) project and other projects world-wide. We are exploring a set of innovations in execution models, programming models and methods, runtime and operating system software, adaptive scheduling and resource management algorithms, and instrumentation techniques to achieve unprecedented efficiency, scalability, and programmability in the context of billion-way parallelism. A number of applications have been implemented to drive system development and quantitative evaluation of the HPX system implementation details and operational efficiencies and scalabilities.

## Applications and Characteristic Behavior

- LULESH:** (Livermore Unstructured Lagrangian Explicit Shock Hydrodynamics) For details see Deep Dive to right.
- Mini-ghost:** A mini-app for exploring boundary exchange strategies using stencil computations in scientific parallel computing. Implemented by decomposing the spatial domain, inducing a "halo exchange" of process-owned boundary data.
- N-Body Code:** An event driven constraint based execution model using the Barnes-Hut algorithm where the particles are grouped by a hierarchy of cube structures using a recursive algorithm. It uses an adaptive octree data structure to compute center of mass and force on each of the cubes with resultant  $O(N \log N)$  computational complexity making use of LibGeoDecomp, an auto-parallelizing library.
- PIC:** Two 3D particle-in-cell (PIC) codes – GTC and PICAR. The gyrokinetic toroidal code (GTC) was developed to study turbulent transport in magnetic confinement fusion plasmas. It models the interactions between fields and particles by solving the 5D gyro-averaged kinetic equation coupled to the Poisson equation. PICAR is a mini-app with the key functionalities of PIC accelerator codes, including a Maxwell solver using an arbitrary order finite-difference scheme (staggered/centered), a particle pusher using the Boris algorithm, and an energy conserving field gathering with high order particle shape factors.
- minTri:** A newly developed triangle enumeration-based data analytics minTri. minTri mimics the computation requirements of an important set of data science applications, not well represented by traditional graph search benchmarks such as Graph500. An asynchronous HPX-based approach enables our linear algebra-based implementation of minTri to be significantly more memory efficient, allowing us to process much larger graphs.
- CMA:** (Climate Mini-App) For details see Deep Dive to right.
- Kernels:** Various computational kernels, such as matrix transpose and fast multipole algorithms, which are used to explore features of HPX and compare to other approaches.

### LULESH (Deep Dive)

LULESH solves the Sedov blast wave problem in three dimensions, the problem is spherically-symmetric and the code solves the problem in a parallelized region. In the figure, symmetric boundary conditions are imposed on the colored faces such that the normal components of the velocities are always zero. Free boundary conditions are imposed on the remaining boundaries.



The LULESH algorithm is implemented as a hexahedral mesh-based code with two centering. Element centering stores thermodynamic variables such as energy and pressure. Node centering stores kinematical values such as positions and velocities. The simulation is run via time integration using a Lagrange looping algorithm. There are three main computational phases within each time step: advance node quantities, advance element quantities, and calculate time constraints. There are three communication patterns, each regular, static, and uniform: face adjacent, 26 neighbor, and 13 neighbor communications, illustrated below.



[1] LULESH is available from: <http://webpages.itd.gov/~huckshy>

[2] Karlheinz G. Kuster, J. Cheminot, H. Cohen, J. D. V. et al. Exploring Traditional and Emerging Parallel Programming Models using a Policy Application. In Proc. of the 27th IEEE International Parallel and Distributed Processing Symposium (IPDPS), 2013.

### CMA (Deep Dive)

The climate mini-app (CMA) models the performance profile of an atmospheric "dynamic core" (dcore) for non-hydrostatic flows. The code uses a conservative finite-volume discretization on an adaptively-refined cubed-sphere grid. An implicit-explicit (IMEX) time integrator combines a vertical implicit operator (which is FLOW-bound) with a horizontal explicit operator (which is bandwidth-bound). There are three major sources of load imbalance in the code:

- The number of iterations required to perform the non-linear vertical solves will vary across the grid.
- Exchanges across the boundaries of the six panels of the cubed sphere are more computationally expensive than intra-panel exchanges.
- The adaptively refined mesh will add and removed refined regions as the simulation evolves.

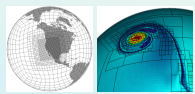
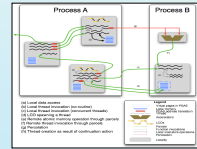


Figure: Left image shows an example of an adaptively refined cubed-sphere grid used in climate codes. Right image shows velocity dynamics for a climate test problem with AMR.

The mini app is implemented using the Chombo adaptive mesh refinement (AMR) framework and has both an MPI+OMP and HPX backends. The mini-app is being used to explore performance on multi-core architectures (e.g. Xeon Phi) and to explore the benefits of using HPX for finite-volume AMR codes to combat dynamic load imbalance.

## Background Info

### Parallel Execution Model

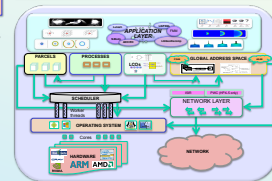


- Lightweight multi-threading
  - Divides work into smaller tasks
  - Increases concurrency
- Message-driven computation
  - Move work to data
  - Keep work local, stage blocking
- Constraint-based synchronization
  - Declarative criteria for work
  - Event driven
  - Eliminates global barriers
- Data-directed execution
  - Merger of flow control and data structure
- Shared name space
  - Global address space
  - Simplifies random gathers

### High Performance Parallel (HPX)

- The HPX runtime system relies the ParallelX execution model to support large-scale irregular applications:
  - Localities
  - Active Global Address Space (AGAS)
  - ParallelX Processes
  - Complexes (ParallelX Threads and Thread Management)
  - Parent/Transport and Parent Management
  - Local Control Objects (LCOs)
  - Size between the application and OS
- Portable interface: C++11/14 (HPX-3 only), XPI
- Comprehensive suite of parallel C++ algorithms (HPX-2 only)
- Automatic distributed garbage collection in AGAS (HPX-3 only)
- Flexible set of execution and scheduling policies
- Performance counter framework (HPX-2 only)

### HPX Architecture



HPX avoids the use of locks and/or barriers in parallel computation through the use of LCOs, which are lightweight synchronization objects used by threads in a control mechanism. Reads and writes on LCOs are globally atomic and require no other synchronization mechanism.

HPX has optimized transports built on top of Photon (HPX-3 Only) and other communication libraries

- Two-sided hand-free transport (ISR)
- Pre-points access to reduce probe overhead
- One-sided Pair-Walks-Command-completion (PWC)
- Local remote modifications for RDMA operations
- RDMA communication optimizations using Photon: turn large push into gets, buffer coalescing

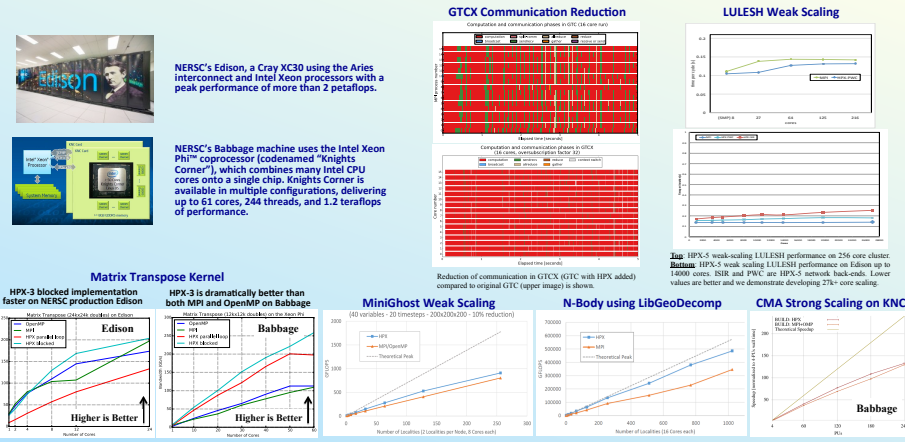
### Memory Models and Transport



HPX supports three memory models: Symmetric Multi-Processing (SMP, no global memory), a Partitioned Global Address Space (PGAS), an Active Global Address Space (AGAS).

- [1] Thomas Sterling, Daniel Kogler, Matthew Anderson, and Marco Brondino. SLEWBER: A performance model for fractal computing. Supercomputing Frontiers and Innovations, 1:42–57, September 2014.
- [2] Hartmut Kaiser, Thomas Heller, Bryce Adelstein-Lelbach, Adam Serio, Dierker Fox, XPI: A Task-Based Programming Model in a Global Address Space. PDCAS 2014: The 16th International Conference on Partitioned Global Address Space Programming Models (2014).

## Results Showing Benefits of HPX



## Performance Adaptation, Legacy Applications, and Summary

### APEX : Performance Adaptation

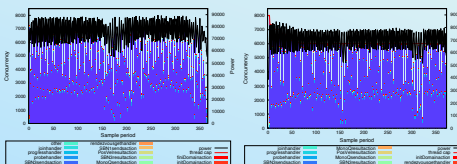
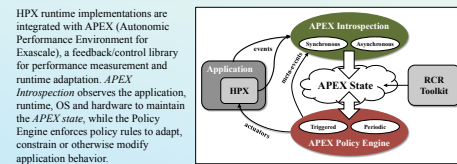


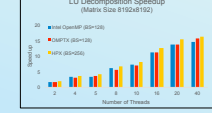
Figure: Concurrency views of LULESH (HPX-5), as observed and adapted by APEX. The left image is an unmodified execution, while the right image is a runtime power-capped (220W per-node) execution, with equal execution times. (8000 subdomains, 643 elements per subdomain on 8016 cores of Edison, 334 nodes, 24 cores per node). Concurrency throttling by APEX resulted in 12.3% energy savings with no performance change.

### Legacy Application Support

OMPITX is an HPX implementation of the Intel OpenMP runtime, enabling existing OpenMP applications to execute with HPX.



Within a node, performance achieved for OpenMP programs using OMPITX is comparable to using the Intel OpenMP Runtime. For example, below we show speedup for a blocking LU decomposition benchmark using the optimal block size for each implementation. These results were obtained with a dual-socket Ivy Bridge processor at LSU.



### Summary

Exascale programming models and runtime systems are at a critical juncture in development.

Systems based on light-weight tasks and data dependence are an excellent method for extracting parallelism and achieving performance.

HPX is emerging as an important new path with support from US Department of Energy, the National Science Foundation, the Bavarian Research Foundation, and the European Horizon 2020 Programme.

Application performance of HPX codes on very recent architectures, including current and prototypical next-generation Cray-Intel machines, is very good.

For some applications, the performance using HPX is significantly better than standard MPI + OpenMP implementations.

Performance adaptation using APEX provides significant energy savings with no performance change.

We have shown that legacy applications using OpenMP can run under the HPX runtime system effectively.



Funding by DOE Office of Science through grants: DE-SC0008714, DE-SC0008809, DE-AC04-94AL85000, DE-SC0008596, DE-SC0008638, and DE-AC02-05CH11231, by National Science Foundation through grants: CNS-1147470, AST-1240655, CCF-1166602, IIS-1447831, and ACI-1339782, and by Friedrich-Alexander-University Erlangen-Nuremberg through grant: H2020-EU.1.2.2. 671603.

