Extreme Scale Software Stack Program
Summary of Impact on Science

DOE Program Manager: Sonia R. Sachs

DOE Research Director: William Harrod

September 2016

DEGAS Project: Impact of advances in data Structures and Runtime Support for Irregular Data-Intensive Applications

Speedups

- Distributed hash table
  - Applications: HipMer (genomics)

720x

- Irregular data exchange
  - Applications: AMR, HPGMG

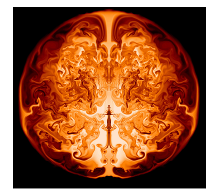

1.2x

- Irregular global matrix update
  - Applications: NWChem, seismic tomography

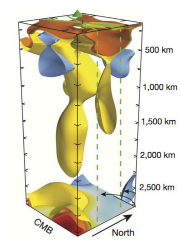

6x

- Distributed work queue
  - Applications: NWChem, Hartree-Fock

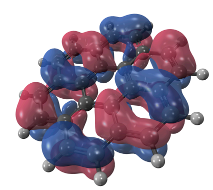

1.2x

- Dynamic task graph
  - Applications: Sparse symmetric matrix solver

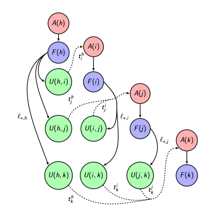

2x

D-TEC Project: Impact of advances on DSL technologies, compilers and runtime systems

- AMR Shift Calculus DSL with ROSE/Polyopt
  - 7.9x for 3D 125 pts stencil
  - 7.3x for 2D 81 pts stencil

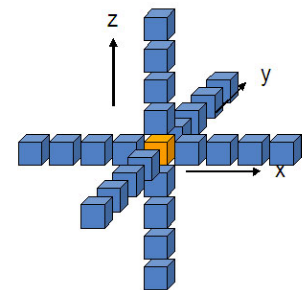

- Halide DSL
  - 4.25x for MiniGMG
  - 1.8x on GPU for HPGMG

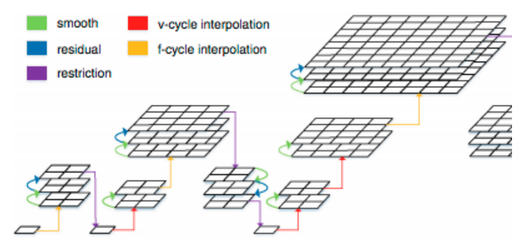

- Bamboo
  - 1.27x in 32K size for algebraic multigrid
  - 1.29x with 96K grid cells for 3D stencil
- LULESH with X10
  - 1.12x better performance
  - 40% fewer lines of code

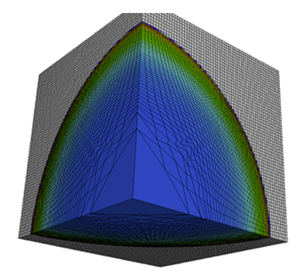

- Global-View Resilience (GVR)
  - 85% parallel efficiency on 1K processes with less than 2% code change

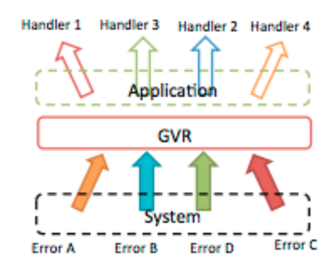

Traleika Glacier Project: The Open Community Runtime Software Suite and its Impact on Applications

- Applications: Full set of DOE mini-apps and the full app atmospheric circulation, Tempest:
  - Tempest (!!)
  - SCF from NWCHEM
  - CoMD
  - HPCC and HPCG kernels
  - Lulesh (multiple versions)
  - miniAMR
  - HPGMG
  - **Genomics Smith-Waterman**
- Full OCR API supported on real hardware and is exploited by these tool chains
  - C library, C++library
  - CnC on OCR, Hierarchically Tiled Arrays (HTA) on OCR
  - Compiler generation of OCR calls (R-Stream)
  - Habanero-C language on OCR

Open
Community
Runtime

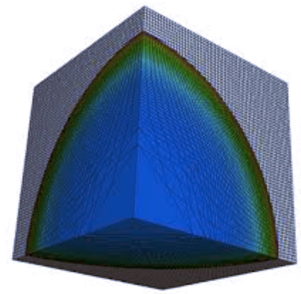

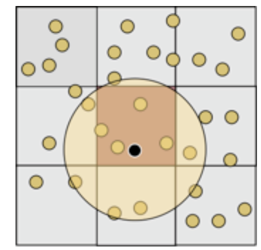

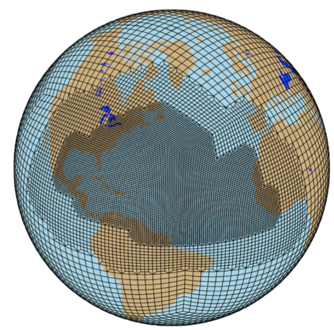

Please see:
https://xstack.exascale-tech.com/wiki/index.php/Main_Page#Traleika_Glacier_Research_Products and https://xstack.exascale-tech.com/wiki/index.php/Traleika_Glacier_Software_Releases
for details, products and research successes

XPRESS Project: impact of exascale runtime support (HPX)

		Comparisons/Results
• N-body Simulation	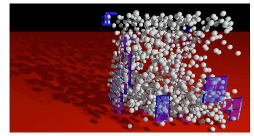	**1.4x** over MPI (16,384 cores)
• Mini-Ghost: Boundary Exchange Mini-app		**1.13x** over MPI+OpenMP (1024 cores)
• Kernels: Stream benchmark Matrix transposition		**1.4x** over OpenMP **2.5x** over MPI+OpenMP
• Distributed GPU work	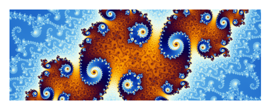	**1.5x** over native CUDA on 16 GPUs
• Lulesh: Shock Hydrodynamics	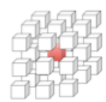	**1.2x** over MPI on Cori (128 cores) **Same** as MPI on Cori (4k cores)
• DSEL and MTL for HPX • Same Portable code GPU / CPU	MTL4 DSL for linear algebra through DOE NNSA DE-NA0002377 (PSAAP2)	**Same** as MPI (256 cores)
• Photon: Integrated Communication Library	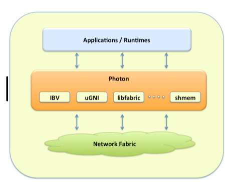	**1.34x** for 16 byte puts **1.37x** for 16 byte gets over MPI-3 one-sided

PIPER Project: Impact of PIPER Technology on …

… Application Performance

- Active Harmony+Caliper auto-tuning **converged 40% faster** on **Lulesh**
- Blame Data Centric profiling of **a high-level parallel language** led to **38-56% speedup**
- Network bottleneck analysis on **NWChem** led to **20% speedup**

… Application Development

- Caliper has been integrated into the hypre solver library and LLNL IC codes
- New tools interfaces have started to become widely used, providing users with critical information during development

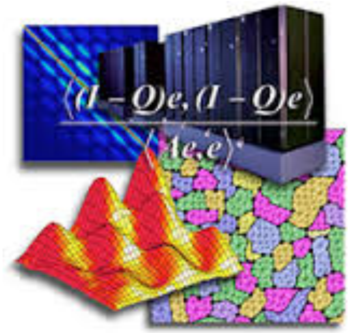

… The ECP/ATDM SW Stack

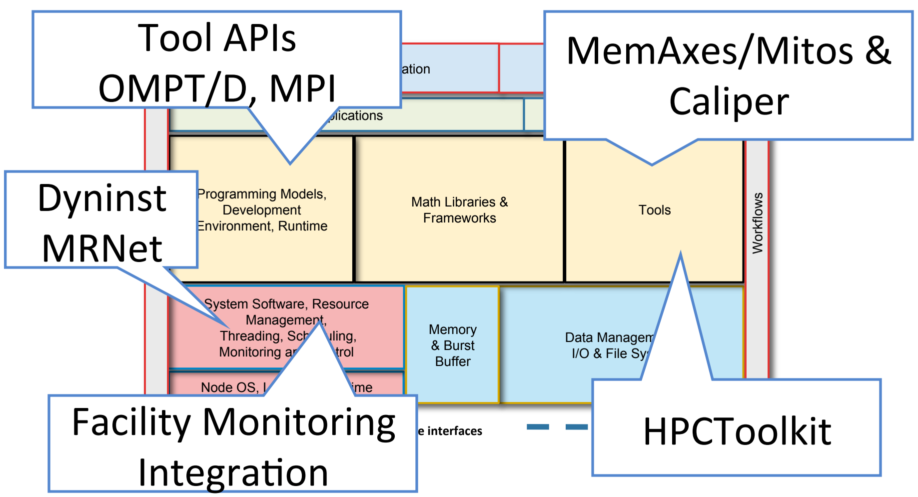

… Advances in R&D

- Tool integration between PIPER components
- The PALM Modeling Toolkit
- Networks and Contention Analysis

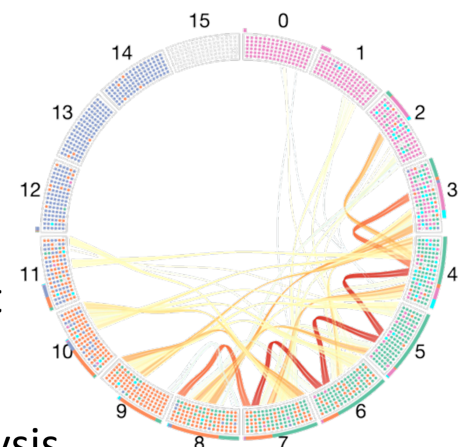

X-TUNE Project: Autotuning for Exascale

Domain-specific and standard compiler transformations combined with autotuning achieve high performance and improve programmer productivity.

Motifs

STENCILS & GEOMETRIC MULTIGRID

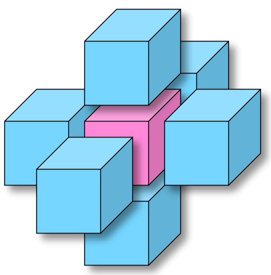

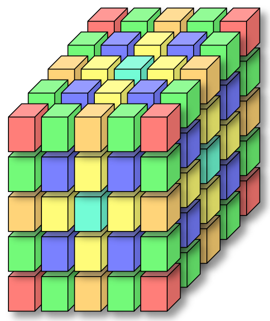

miniGMG benchmark, proxies the MG solves in BoxLib/Chombo codes (ExACT)

Impact

Solver, 7pt GSRB variable coefficient stencil, & 125pt Jacobi constant coefficient stencil

Speedup over reference (CPU): **3x**

Reduction in lines of code: **>10x**

Performance portability: Outperforms manually-tuned code on CPU and GPU

High performance: Near roofline model bound

Scalability: Demonstrated on 1K nodes

TENSOR CONTRACTION & SPECTRAL ELEMENT

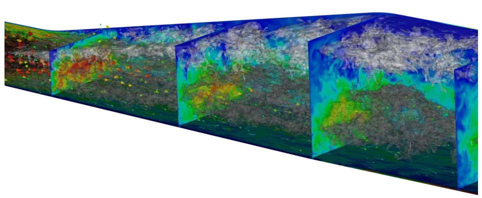

Nekbone benchmark, proxies Nek5000 (CESAR)

Other relevant application: NWCHEM

Speedup (GPU) over OpenMP (CPU): **1.5x**

Speedup over tuned OpenACC: **2.9x**

Reduction in lines of code : **>100x**

Fully automated: Mathematical formula to high-performance CUDA

Performance & productivity: Autotuning essential even for OpenACC code

Corvette Project: Dynamic Analysis for Program Verification and Optimization

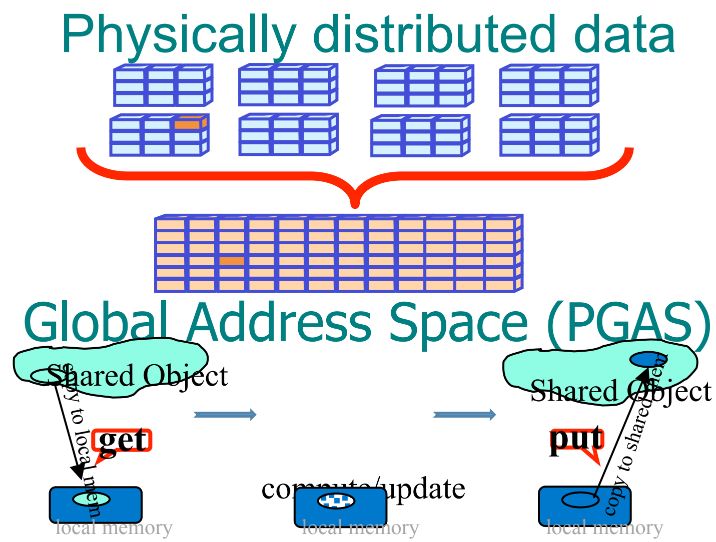

- Scalable data race detector for PGAS languages
  - 50% overhead at 8K cores , 200X faster than commercial tools
- Eliminating redundant synchronization
  - NWChem -> 14% speedup at 2K cores
- Exploiting performance variability for energy optimizations in dynamic apps
  - NWChem – 20% energy savings at 2K cores
- Dynamic program analysis for communication optimizations
  - HPGMG - 65% less time spent in communication
- Floating point reproducibility
  - ReproBLAS - 1.2x to 3.2x slowdown vs. fastest non-reproducible code
- Floating point precision tuning
  - lowered precision in Gnu Scientific Library, up 40% speedup

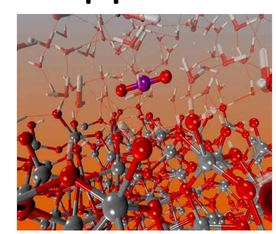

credit:nwchem-sw.org

HPGMG

SLEEC Project: Semantics Rich Libraries for Effective Exascale Computations

Speedups

- Optimized scheduling for recursive domain decomposition
  - Applications: Computational mechanics, peridynamics
  - Publications: ICS 2013

2.5x

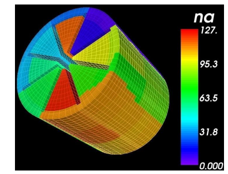

- Domain-aware partitioning strategies
  - Applications: Recursive coupling applications
  - Publications: IJNME 2016 (submitted)

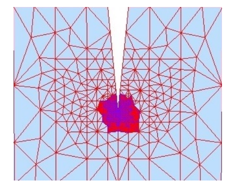

11x

- Semantics-aware multi-accelerator offloading
  - Applications: Heterogeneous apps (e.g., Jacobi)
  - Publications: ICS 2013, ICS 2015

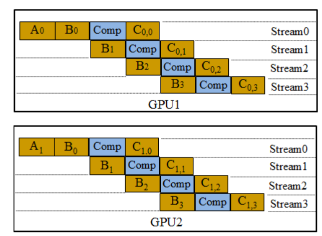

130x

- Semantics-based optimizations for CnC
  - Applications: LULESH
  - Publications: WolfHPC 2015, LCPC 2015

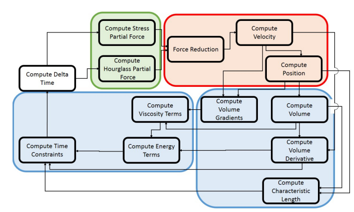

120x

GVR Project: Robust Resilience for High Error Rate Environments/Systems

- Expand ABFT from immediate to Latent and Silent Errors at extreme scale
  - Efficient Versioning and Recovery library
  - High performance, scalable versioning (NVM)
- Deep App Studies & New Recovery Types
  - Monte Carlo: OpenMC, Particle: ddcMD, AMR: Chombo, Iter: PCG/Trilinos.
  - Rollback, Adaptive, and Fwd Approx Recovery
  - 16,384 Rank experiments: Scalable & High Performa
  - Practical: Only Localized Code Change

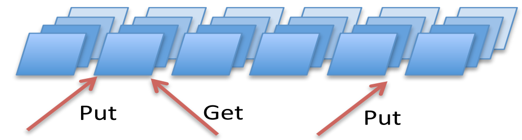

A. 1000's of Fast versions (NVM,SSD)

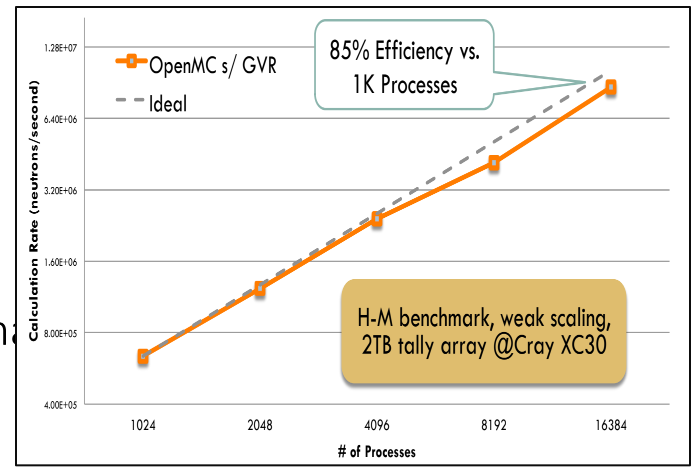

B. Scales well, demonstrated >16K Ranks

C. Practical: < 1% change

Application	% Changed	Application Lines of Code	Leverage Global View
OpenMC	<2%	30 K	Yes
PCG/Trilinos	<1%	300 K	Yes
ddcMD	<0.3%	110 K	Yes
Chombo	<1%	500 K	Yes

http://gvr.cs.uchicago.edu

Andrew A. Chien & Pavan Balaji

DOE/ASCR DE-SC0008603, DE-AC02-06CH11357

THE UNIVERSITY OF CHICAGO Argonne NATIONAL LABORATORY

Vancouver: Improving Programmability of Contemporary Heterogeneous Architectures

- Understanding novel heterogeneous architectures
  - SHOC Benchmarks
  - Application engagement and refactoring
- Developing languages and compilers to facilitate portability
  - OpenARC compiler infrastructure for GPU, Xeon Phi, FPGAs
  - KLAP – CUDA GPU Dynamic parallelism compiler
- Building autotuning frameworks that hide complexity
  - Tanagram – kernel synthesis
- Designing scalable performance analysis and modeling tools
  - Scalable performance tools for heterogeneous systems - Tau
  - Automatically generating performance models - COMPASS
- Deployed open-source tools

ARES: Abstract Representations for the Extreme-Scale Stack

- Create a universal high-level intermediate representation
  - HLIR defined that includes heterogeneous computing and complex memory hierarchy concepts
- Develop prototype frontends
  - OpenACC and OpenMP offload frontends developed and deployed
- Develop prototype optimization engine for HLIR
  - Optimization engine recognizes and optimizes
- Develop back-end compilation, based on LLVM
  - HLIR lowered to LLVM IR, gaining benefits of LLVM infrastructure
- Demonstrate HLIR benefits on application examples
  - IMPACC, NVL-C, OpenACC2FPGA, FITL, etc.