


*Exceptional service in the national interest*



# SST and ExMatEx Update

S.D. Hammond, A.F. Rodrigues, S.M. Kelly, J.Q. Wilson and J.P. Vandyke  
Scalable Computer Architectures, Sandia National Laboratories



# SST and Co-Design Update

- Improvements in the SST Toolkit and plans for 2013
- New tools and approaches for 2013
- Overview of co-design analysis which is underway
- ExMatEx objectives for forthcoming year

# IMPROVEMENTS IN SST

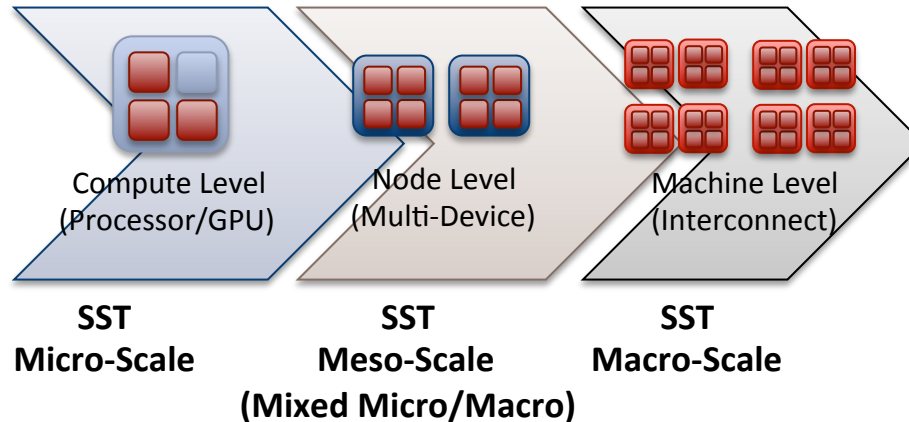
# Overview of Changes in SST

- Sandia has formed a “productization” group for SST
  - Software-engineering expertise – testing, documentation, quality assurance *etc*
    - Now made three supported releases to the community
    - Seen increased adoption of SST based on questions + answers, downloads, *etc*
  - On schedule for a significant release in April 2013
  
- Validation activities have begun
  - Driving SST components to improved levels of accuracy
  - Enhancing our understanding of simulation technology in general
  - Active collaboration with Sandia’s ASC V+V team
    - Driving verification, validation and calibration activities

# Validation Activities

- Accuracy of SST predictions is challenging
  - Usually more accurate than analytic models for complex workloads
    - Captures interaction/congestion across multiple components
  - Significant cost (runtime) associated with performing studies and analysis – means refining models is very time consuming
  
- Studies now underway to actively verify important SST components
  - Particular focus on memory systems
    - DRAM STREAM models now verify successfully to ~85-90%
    - Memory latency models verify to ~85-95%
  - Driven by vendor questions

# Integration of Macro/Micro



- First integration of SST/Macro with the Micro code base
  - Enables Macro-level modes to be executed within Micro
  - Swap-in of Micro-level network components
  - Prototypes of OpenSHMEM applications running through Macro
- Potential for integration of UPC through MPI layers into Macro (being considered for FY13 work)
- Well developed studies underway for network topology analysis

# Threading in SST/Micro

- Initial work to allow simple pthread implementations to run within Micro
  - Turns out this has been extremely challenging
    - Getting memory buses to lock correctly and give coherent answers is hard
- Does not support pthread condition updates
  - Means direct compilation of OpenMP will not work
  - Potential for simpler compilation through XOMP but this is future work
- Been able to demonstrate scalability of small mini-applications from Mantevo, STREAM and GUPS

# Cache Models and Memory

- New cache modeling components are in final stages of development
  - Support locked buses and coherency (currently on MESI snoop, future should enable directory-based protocols)
  - Enables higher scalability for processor design studies
    - Utilize existing core model for compute
    - Connect many cores through new cache designs
- Expecting that this will be a significant advanced in studying cache designs and higher core counts
  - Internal studies completed in 2012 shows we can produce reasonable estimates of memory bandwidth for Intel's Xeon Phi product line

# BROADENING TOOLS AND APPROACHES

# Broadening Tools Base

- SST was always designed to:
  - Provide flexible component integration (you can build your dream machine through a shopping list of components)
  - Range from cycle-accurate components to analytic-style models
    - Compute models today range from cycle-accurate X86-64 to stochastic Niagara/SPARC models
- SST as project is broadening its tools and component base
  - Including PIN-based tools to generate traffic
  - Statistics and code characteristics tools to help parameterize components
- Driven by the V+V style approach of a box of tools
  - Take your position based on a *body of evidence*

# Why?

- Simulations take a long time
  - Ensemble simulations are taking too long to be useful
  - Turnaround time for initial analyses is also taking too long
- Vendor co-design questions don't always *need* simulation
  - Although tempting to believe simulation answers all of these questions
- The real problem for starting co-design is that we have a very poor understanding of our algorithms
  - Vendors are initially asking for characterization *not* for simulation results

# What does this mean for ExMatEx?


- Broader range of possible analysis
  - Emulation and simulation, complements GREMLINs and ASPEN
- More directed simulation
  - Use emulation to act as a clearing house for some of the analyses
  - Focus our simulations towards studies of real interest
- Clearer path to developing simulation models
  - Lower levels of ambiguity in the model development process

# ANSWERING CO-DESIGN QUESTIONS

# Memory and Cache Sizing


- Vendors are asking many questions about caches, scratch pads and memory design
  
- Tradeoffs associated with:
  - Cache / Scratchpad capacity
  - Access arrangement – cache line, byte-addressable (texture cache)
  - Levels of cache?
  
- Research questions:
  - Can applications take advantage of these features?
  - How can these be integrated into an application?
  - Effect of performance of these architectures?

# LULESH “Touched” Memory



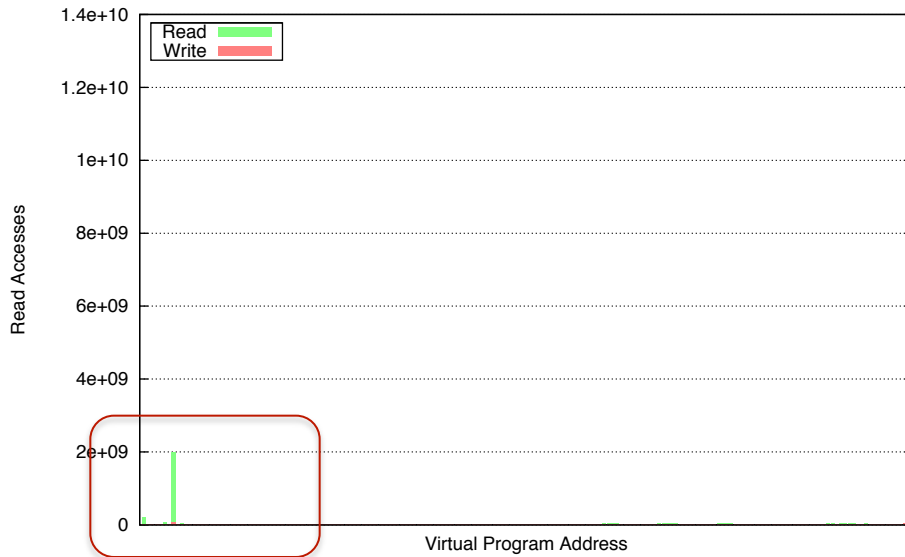
Complete Application Profile

Lower peak

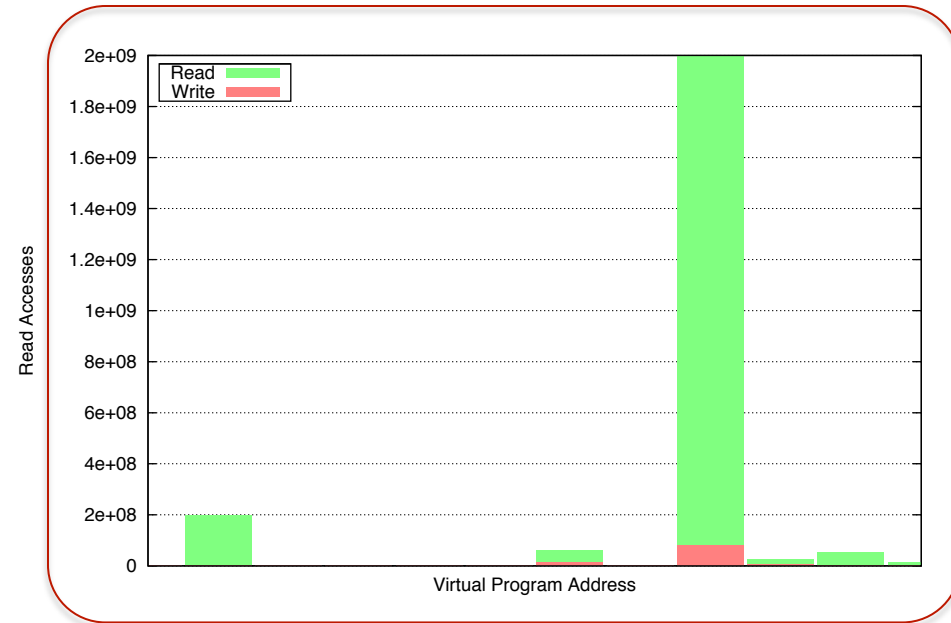


Stack Access Removed

# CoMD “Touched” Memory



Complete Application Profile



Stack Access Removed

CoMD, 8K Atoms (Small), Serial Execution, Touch-Address binning performed at 32kB

# “Touched” Memory Analysis

- Ability to bin memory accesses by read or write into address bins
  - User defined memory bins – everything from page-level to cache-line
  - Includes stack and heap
  - Potential to include instructions (if really needed)
  
- Interesting to look at access type by bin
  - Tells you what you should put into cache and (slower) NVRAM?
  - Tells you approximate size of caches, at least of level-1?
  - Can you partition based on read/write activities?
    - Particularly important in a world where performance of writes may be very slow


# “Touched” Memory Analysis

- Tools part of on-going work with SST and various co-design projects (including ExaCT and ExMatEx)
  - Support from vendors (some already using the analysis output)
- Tools run on unmodified binaries
  - Working for LULESH, CoMD, ASPA, FFT-5 and VPFFT++
- Useful to product estimates of simulations of interest for cycle-accurate SST simulator

# Cache Emulation


- Utilizes simple emulation of multi-level cache
  - Simple cache parameters – size, associativity, cache line width *etc*
  - Allows for arbitrary sizes, hierarchy and arrangement
  - One of the most common questions from vendors
  - Useful for NVRAM, cache sizing and future memory models
- Takes a *long* time to emulate cache behavior
- No timing – that's when we move to SST cycle-timing models

# CoMD Cache Emulation



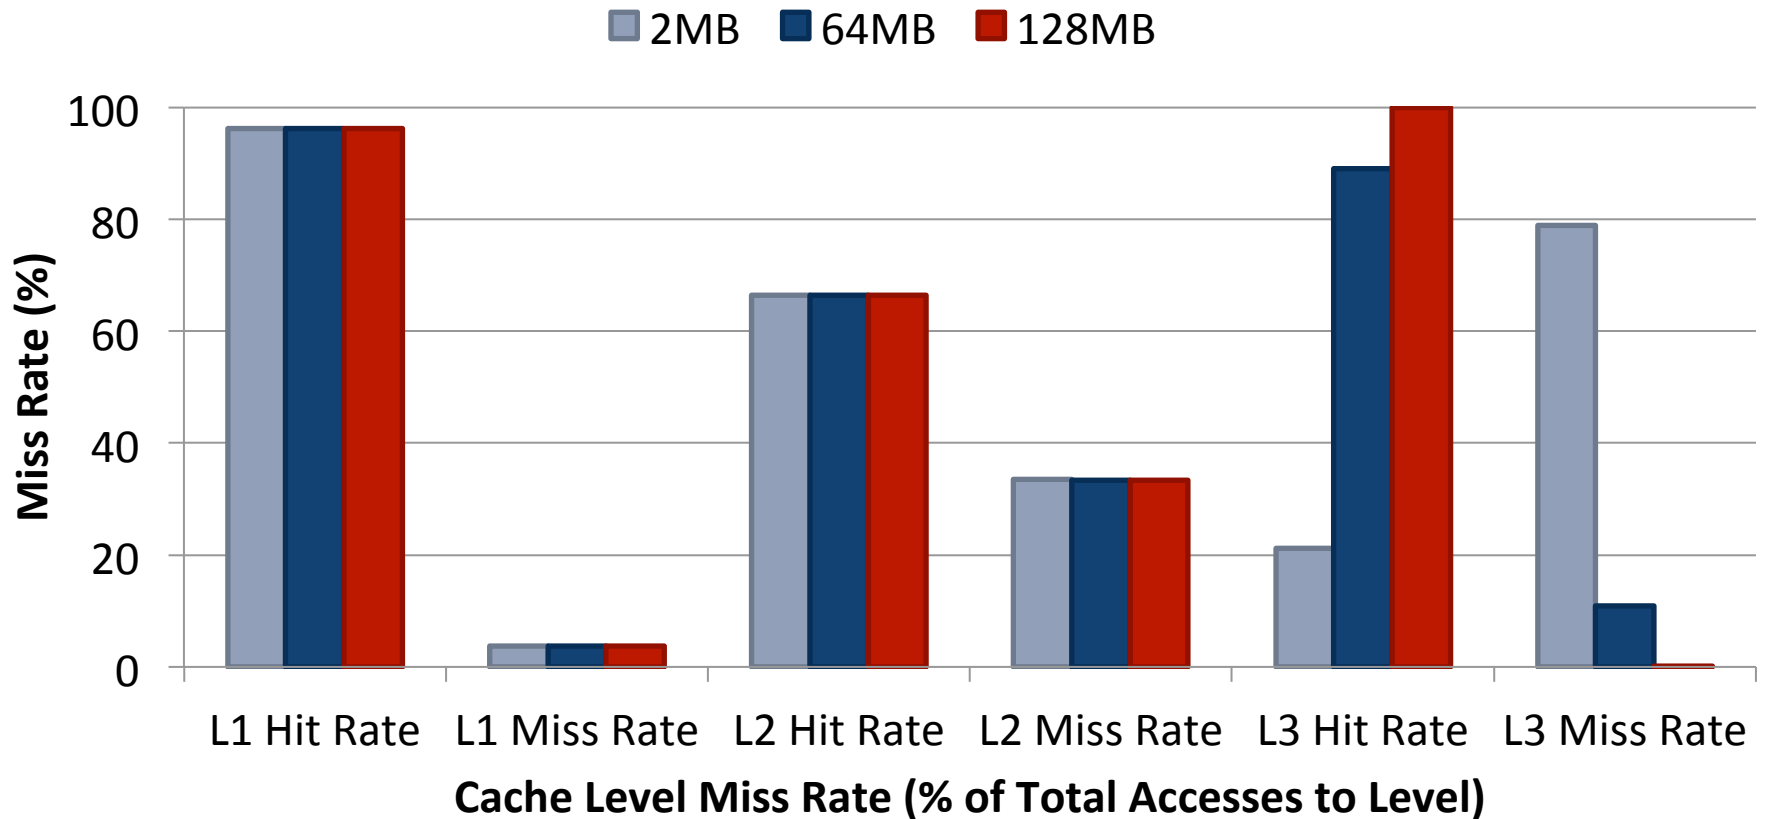
Model is for a standard "8k" input deck, assumes Sandy Bridge split L1I/D, per-core L2 and unified L3, data is for a serial process

# CoMD Cache Emulation



Model is for a standard "262k" input deck, assumes Sandy Bridge split L1I/D, per-core L2 and unified L3, data is for a serial process

# LULESH Cache Emulation



Model is for a standard 50 x 50 x 50 mesh, assumes Sandy Bridge split L1I/D, per-core L2 and unified L3, data is for a serial process

Interesting to see similar (although not identical) shapes for cache misses

# OBJECTIVES FOR EXMATEx

# Cross Validation and Analysis

- SST, GREMLINs and ASPEN
  - Independent approaches to analyzing application performance
  - Complementary if you want cross-validation and analysis
- Potential first steps:
  - Investigating the effects of cache size on mini-app performance
    - GREMLINs can “eat” sections of L3
    - ASPEN can predict memory accesses – question about cache misses?
    - SST can simulation varying cache sizes
  - Algorithm changes in LULESH (for threading)
    - ASPEN can predict tradeoff of FLOP/s and memory
    - SST can run changed code

# Macro models of CoMD

- Have an established model of LULESH for Macro-scale simulation
  - Validated well on BG/L and LLNL's Sierra machine
  
- Next step is an equivalent state-machine representation of CoMD
  - Validation exercises will also need to be conducted
  
- “Roofline” model for scale bridging performance?
  - Using LULESH and CoMD Macro models
  - Predict for sections of a machine running various pieces of code

# Micro models for Mini-Apps

- Expecting to have the ability to do significant analysis of memory and cache bandwidths in coming year
  - CoMD, LULESH and VPPFFT run in SST
  - FFT-5 and ASPA currently do not (ASPA should be fixable)
  
- Use emulation capability and full-scale simulation to produce sensitivity to bandwidths and cache designs
  - Answering one of the most common questions from vendors
  - Helpful to feed into Trinity procurements

# FINAL THOUGHTS

# Some thoughts...

- ExMatEx is challenging SST and our capabilities
  - Mini-apps are larger than Mantevo (in most cases)
    - Means we have to focus on performance and scalability
  - Variety of questions are broader
  - Need to have higher levels of validation and confidence
  
- Actively working to address many of these issues
  - Validation activities underway (which are already improving capabilities and accuracy)
  - New components being added to support the questions vendors are asking us
  - Improving range of tools so simulations are more targeted and more insightful and can deliver a pipeline of results

# Some thoughts...

- We are focusing on co-design questions
  - Already demonstrating initial work is useful
    - Vendors are telling us they need information
  - Blending SST with wider tools is helping us to grow our *body of evidence*
  
- Opportunities to cross-validate and analyze with GREMLINs and ASPEN in coming year
  - Identifying studies to support vendors questions
  - Improving all of our tools
  - Greater confidence in the analysis produced by the center



**Sandia  
National  
Laboratories**

*Exceptional service in the national interest*