SAND2008-2290C

System Simulation
Structural Simulation Toolkit

Arun Rodrigues

Scalable Computer Architectures

CIS External Panel Review

April 13-16, 2008

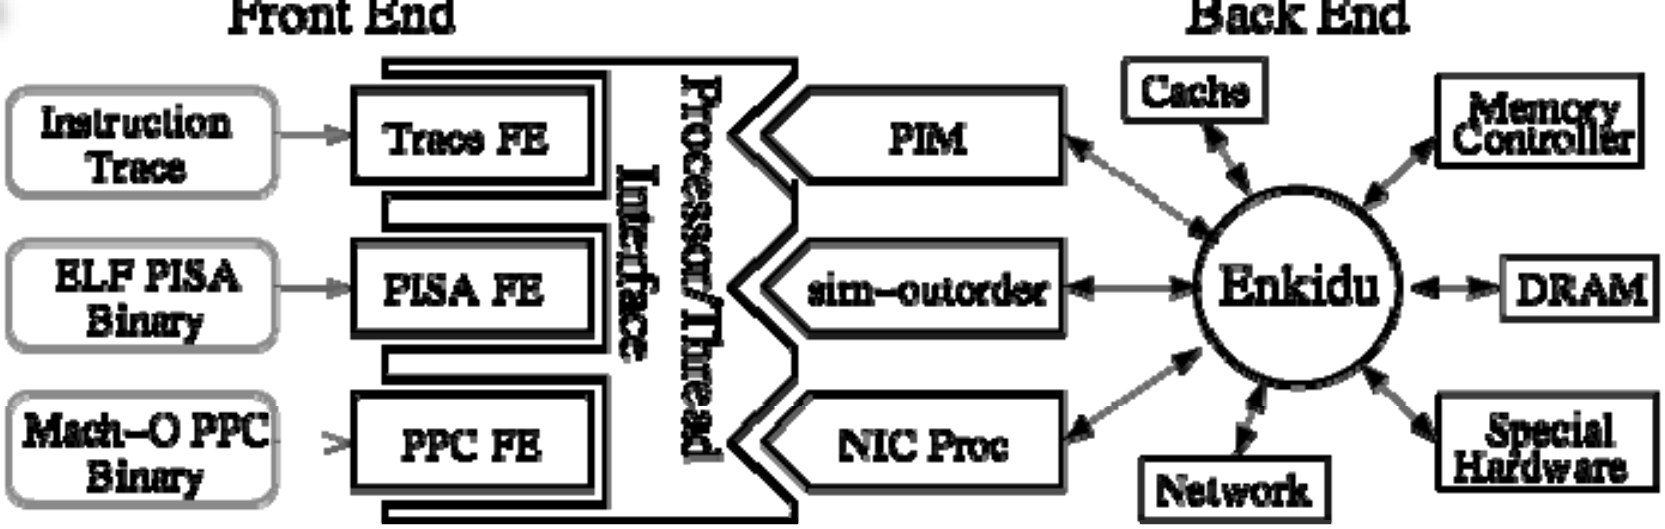

The SST's modular structure allows flexible simulation

The SST was developed to explore systems with novel processors, memories, interconnects and programming mod

Novel Architectures require novel hardware and software

Open Source: http://www.sandia.gov/SST/

SST Features

- **Modular, component-based design**
- **Multiple architectural models**
- **Multi-(processor|threaded|core) simulation**

SST Structure

- **Input drivers (Front-Ends): execution or trace**

```
while(!done) {
  foreach component {
    c->preTic();
  }

  foreach message {
    m->dest->handleParcel(m)
  }
}
```

Enkidu Pseudocode

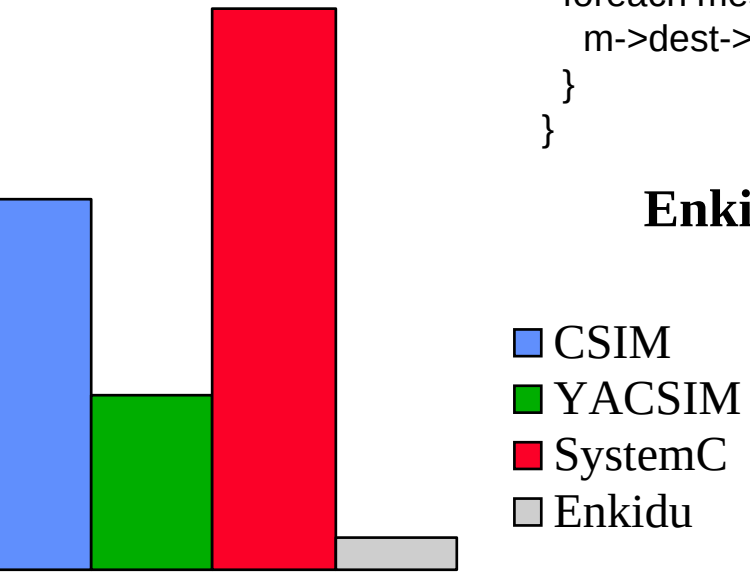

rid DES provides low-overhead

Routine	Simulated	Actual
PUT	0.486	0.592
complete()	0.196	0.154
message()	0.959	1.002
POST	0.477	0.442
message()	1.936	1.686

- **Model of Computation (En**
  - **Component-based hybrid dis event simulation (DES)**
    - **Synchronous / timestep for co internals (analogous to clock)**
    - **DES between components**
  - **Low Overhead (~57 instructi component / cycle)**
- **Hardware component mod**
  - **Cycle-accurate**
  - **C++ behavioral model**
  - **Detail / performance tradeoff**
  - **Several ongoing validation e**

Several Projects

- architecture
- r CacheLine Gather (ICGL)
- on. FU (Wisc./SNL)
- Aggregates
- emory Ops
- cation analysis
- nory Footprint
- ruction Usage
- ork/MPI
- Tradeoffs
- Acceleration
- amming Models
- Compiler work (SNL/Rice)
- allX (LSU/SNL)(FastOS)
- nsactional Memory (ORNL)

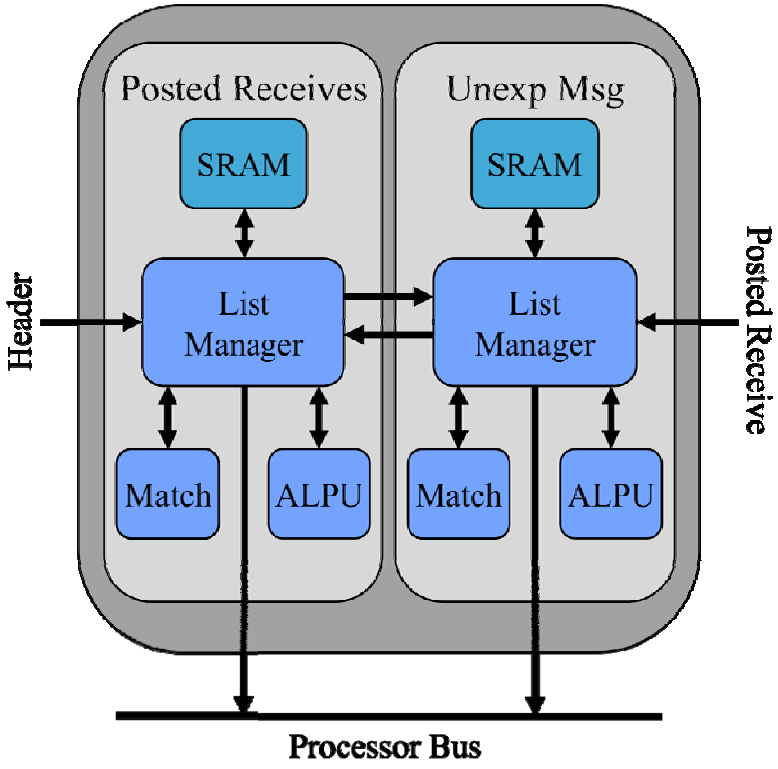

MPI Accelerator Tested with SST

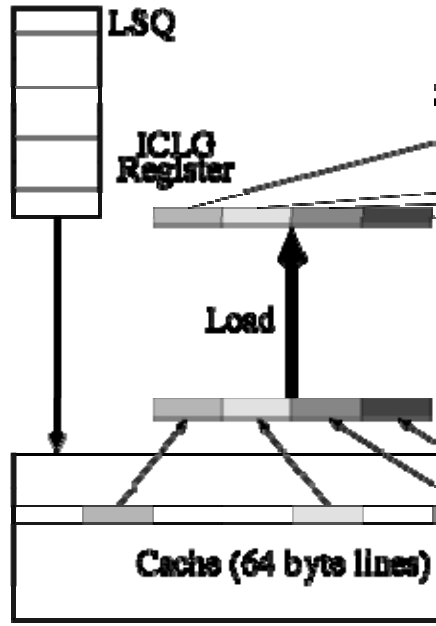

Inter CacheL
Scatter/Gath
Microarchitec

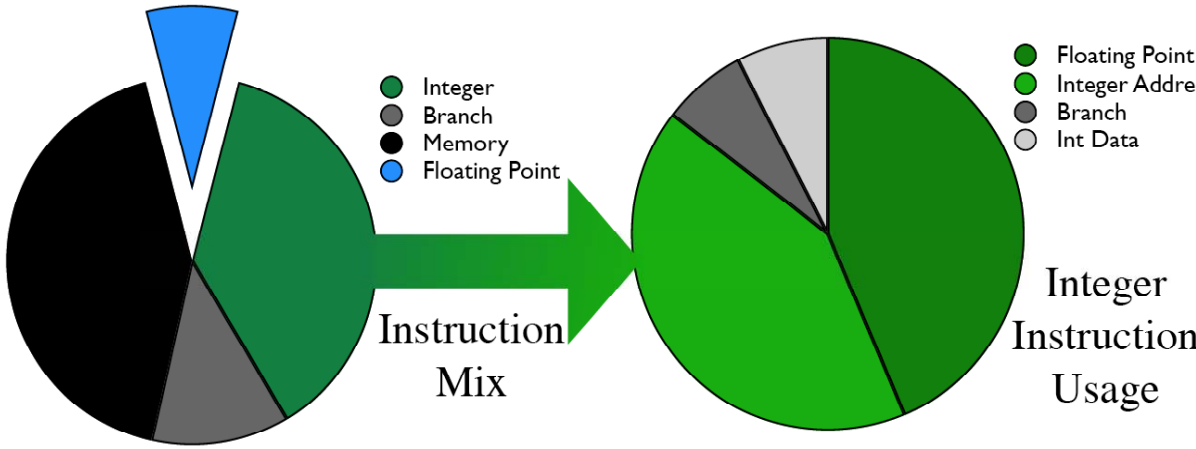

Instruction Mix & Usage for

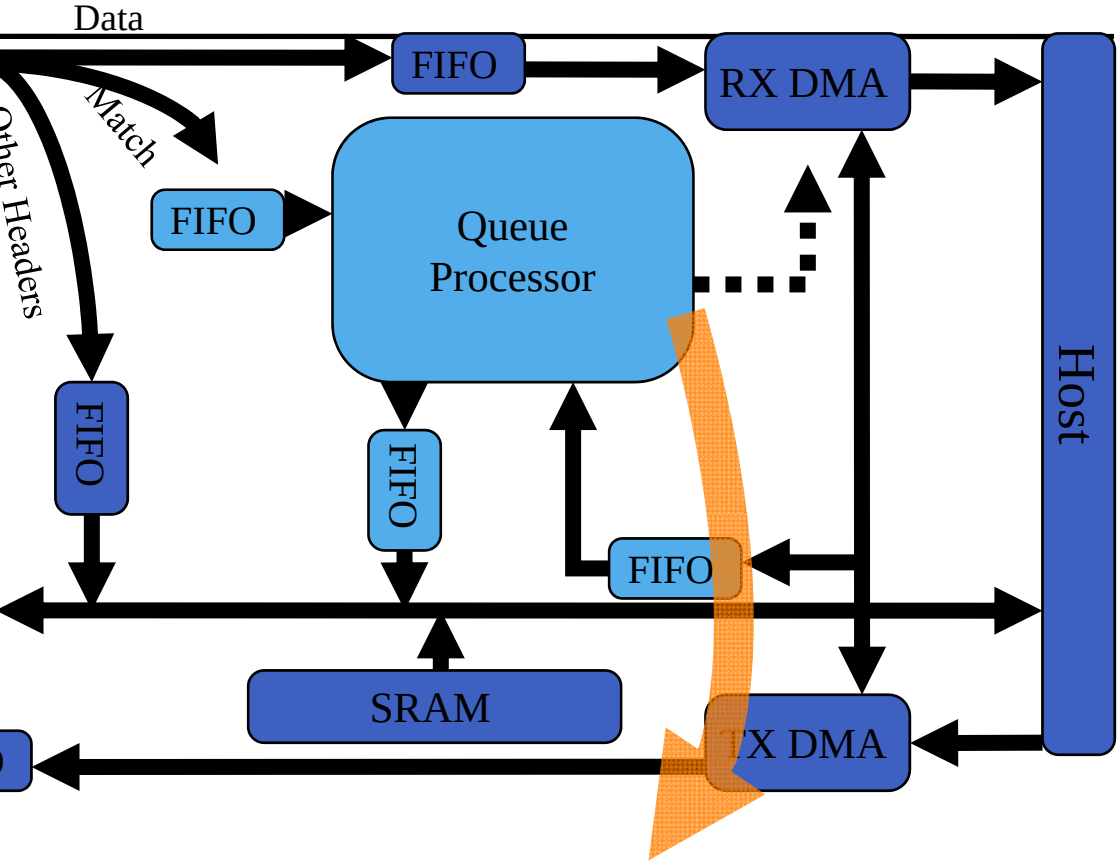

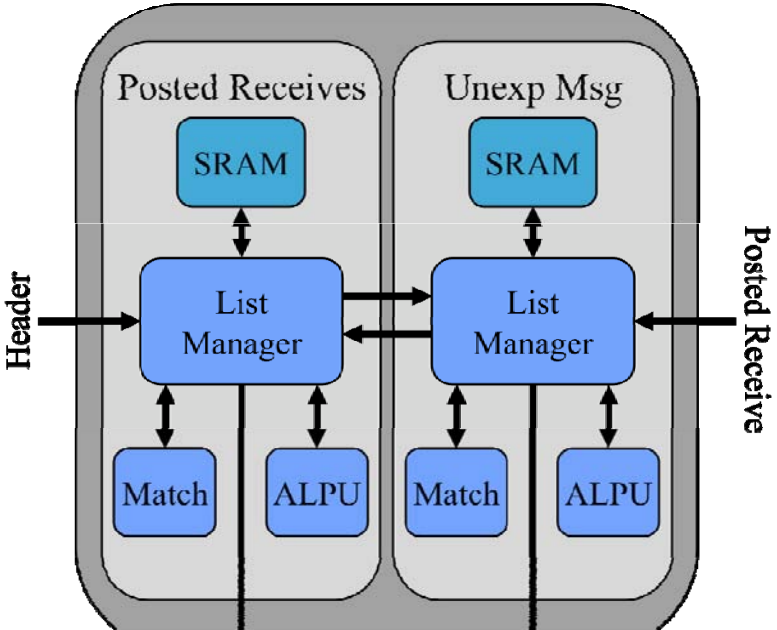

- **Problem: Message rate determines effective banc**
- **SST Analysis: MPI matchi limits performance**
- **Solution: Accelerate list management with hardwa**
- **SST Simulation**
  - **Hardware**
    - **Implement NIC/Router base RedStorm SeaStar**
    - **Implement List Manager, A Match Unit, integrate with N**
      - **"Translated" from FPGA**
    - **Validate against FPGA & R**
  - **Software**
    - **Create baseline offload MP**
    - **Modify MPI to use accelera**
- **Impact: Collaboration w/ I**
  - **Intel licensing**
  - **Keith Underwood "on loan"**

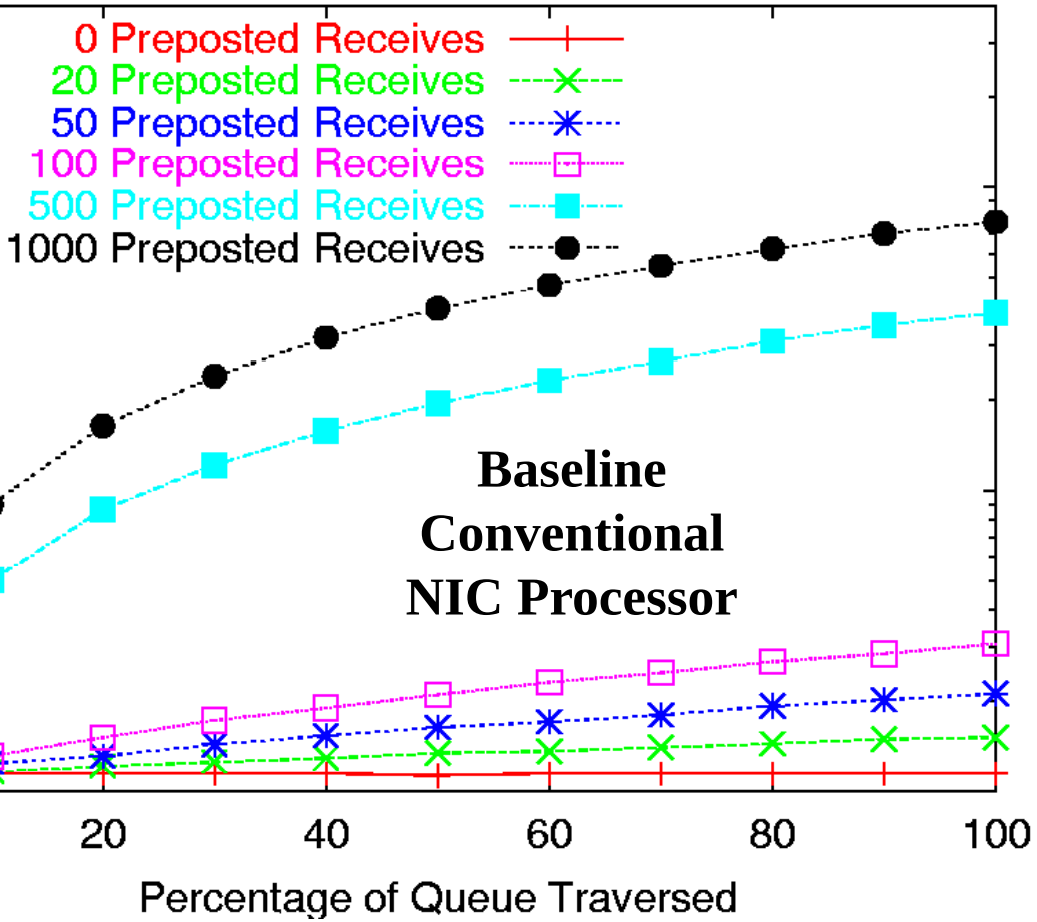

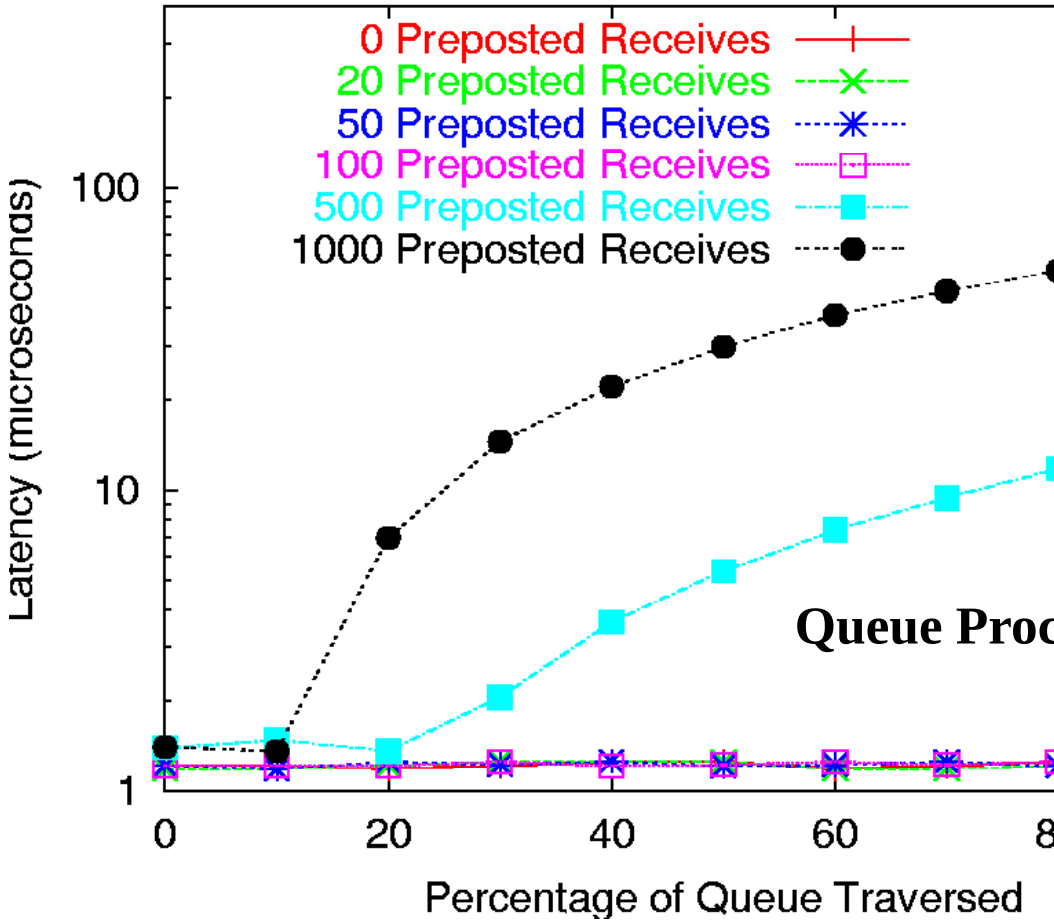

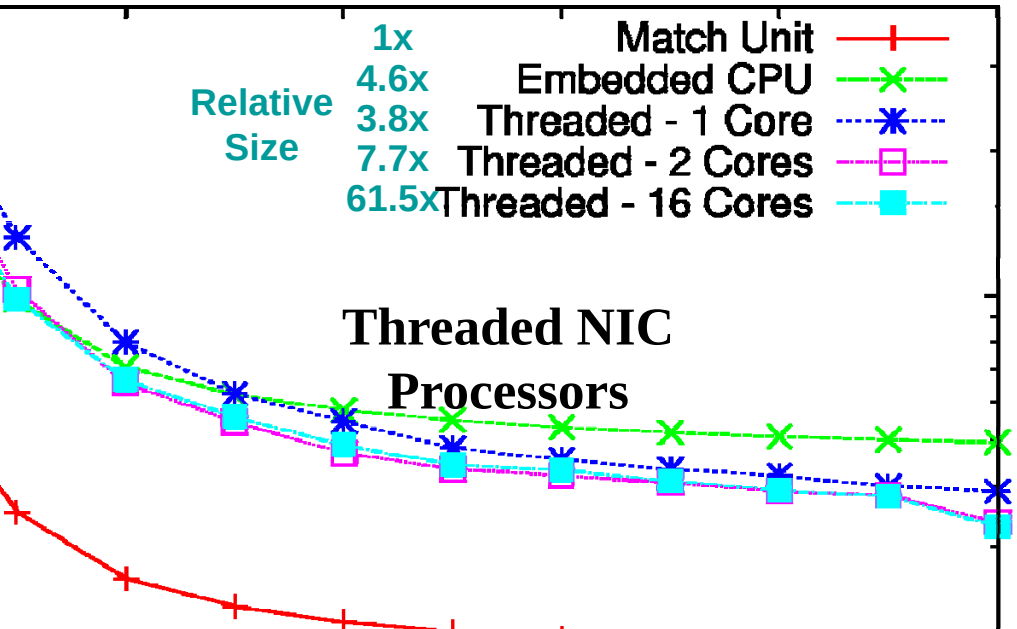

- **Long MPI message que increase effective laten dramatically**
- **Queue processor proce messages more quickly**
- **Queue processor more effective than conventio**

bleScalar & PTLSim (NY
lhamton)

pular

nventional processor only

Network, minimal memory model

g4 (Motorola), Tango (IBM)

gh accuracy

oprietary

ngle model processor only

/MicroLib (Various)

ry detailed

ne grained

ow execution and development
ne

cs / GEMS (Wisc.)

t suited to network study

- LANL/PAL
  - Excellent for scaling, network top studies
  - Requires run on existing proces
  - Can't explore hypothetical proce memory hierarchies
- Srinivasan/Cook Monte Carlo (
  - Extremely fast, accurate
  - In-order processors only

- SST focuses on...
  - "Best of Breed" (don't reinvent)
  - HPC (Network, software stack, ap
  - Parallel (in progress)
  - Abstracted SW/HW interface

ability: Porting / Support / Documentation

w Components

- nproved DRAM Model (Jacob/U. Maryland)
- nproved processor model (PTLSim)

allel execution

- PI extensions to Enkidu
- 000s of simulated nodes on 100s of real nodes
- ready have GTech parallel prototype
  - Parallel DES simplified because of limited communication patterns
  - Allows use of conservative distance-based optimization

tiscale: Tradeoff between precision and ed

- tilize Mantevo models

Simulation

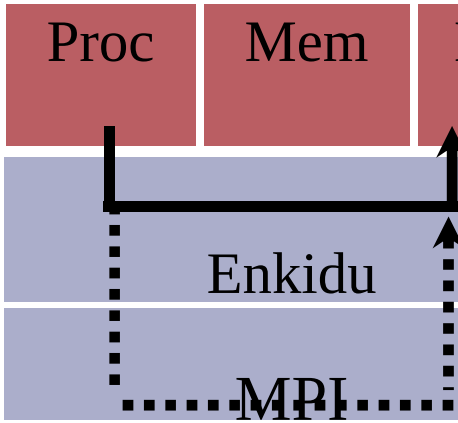

Parallel SST Architec

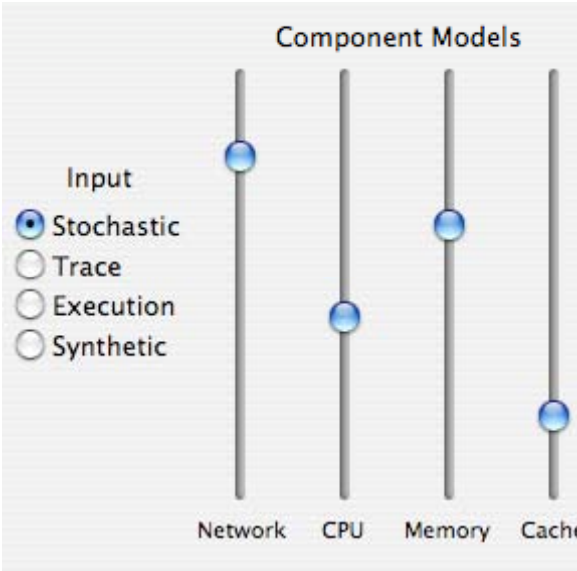

Multiscale SST Param

rchitecture

ereed Papers

chnical Advance

upnow, A. Rodrigues, K. Underwood, and K. Compton. *Scientific Applications vs. SP*
mparison of Program Behavior. In ICS '06, pages 66–74, New York, NY, USA, 2006.

ation analysis

urphy, A. Rodrigues, P. Kogge, and K. Underwood. *The Implications of Working Set A*
upercomputing Memory Hierarchy Design. In International Conference on Supercom
20-22, 2005.

oved understanding of MPI, memory use, instruction use, and application performan

rk/MPI

chnical Advances (2 patents, 4 provisional)

ereed Papers

dation for IAA Network Project

r collaboration with Intel

ssing-in-Memory

ereed Papers

boration with Rice University

dation for IAA Memory Project

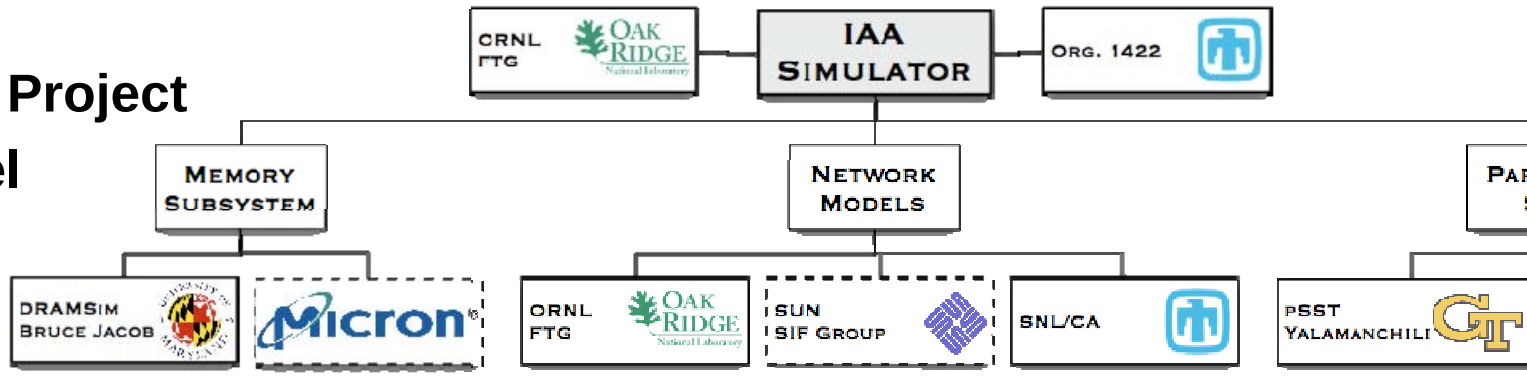

IAA SST Collaborations

odrigues. *Programming Future Architectures: Dusty Decks, Memory Walls, and the S*
, Ph.D. Dissertation, University of Notre Dame, 2006.