

Enhancing an HTM System with Monitoring, Visualization and Analysis Capabilities

**Philipp Kirchhofer¹, Martin Schindewolf¹, Nehir Sonmez², Oriol Arcas²,
Osman S. Unsal², Adrián Cristal², Wolfgang Karl¹**

¹ Karlsruhe Institute of Technology (KIT) ² Barcelona Supercomputing Center (BSC)

Chair for Computer Architecture and Parallel Processing, Institute of Computer Engineering, Karlsruhe Institute of Technology

Transactional Memory simplifies parallel programming

But how to achieve good performance and scalability?

Need for a monitoring infrastructure to

- identify runtime characteristics of an application running on an HTM system
- detect application bottlenecks
- get insight into interaction between application and HTM system

Project Goals

Develop a monitoring infrastructure for the TMbox HTM system using event logs:

- Generate events at run time and save for later processing
- Recreate HTM state offline
- Visualize and analyse saved information


Additional goals:

No probe effect, small hardware overhead, high extensibility, ease of use

Design

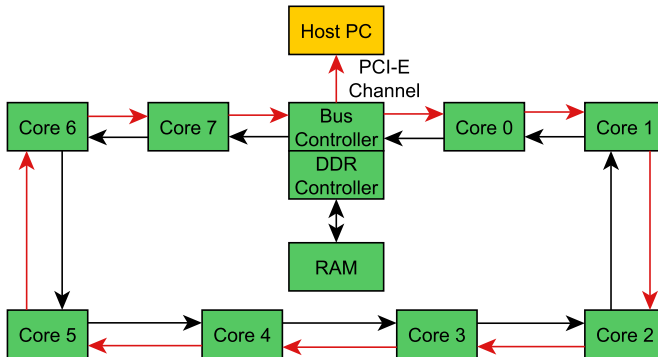
The TMbox system

- Developed at Barcelona Supercomputing Center
- MIPS compatible multi-core system (FPGA based, 16 cores on 1 FPGA)
- Supports
 - STM (TinySTM)
 - HTM (BeeTM)
 - HybridTM (Modified TinySTM)



Design

The TMbox system - Block Diagram

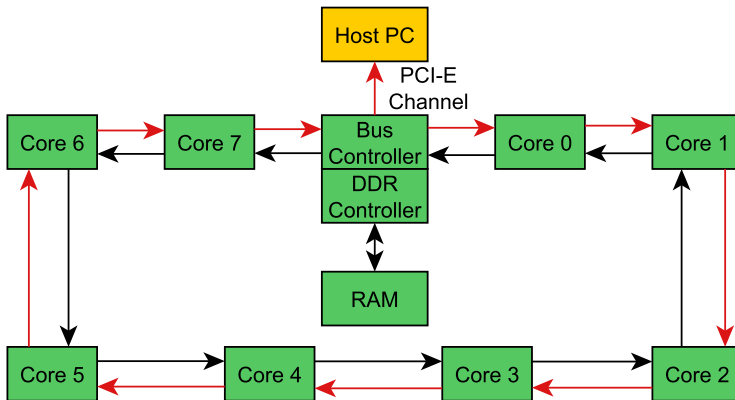


—→ Bus 1: Memory Requests / Responses

—→ Bus 2: Invalidations / Events

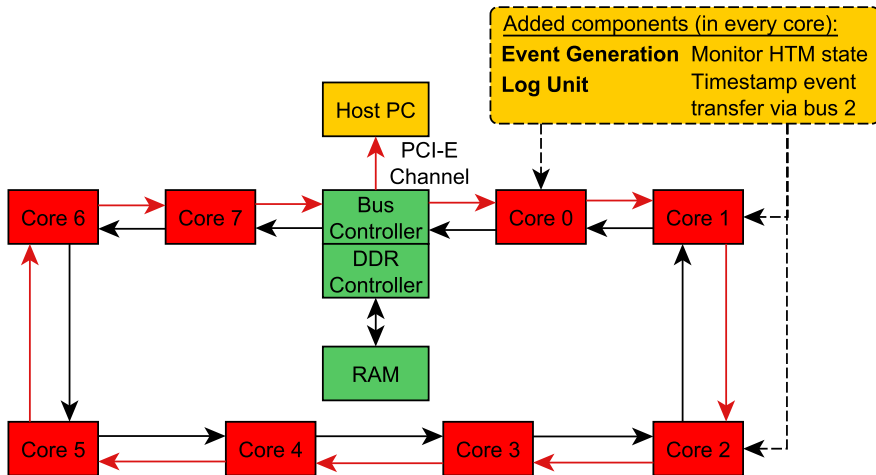
Design

Event Generation, Log Unit, Bus Controller



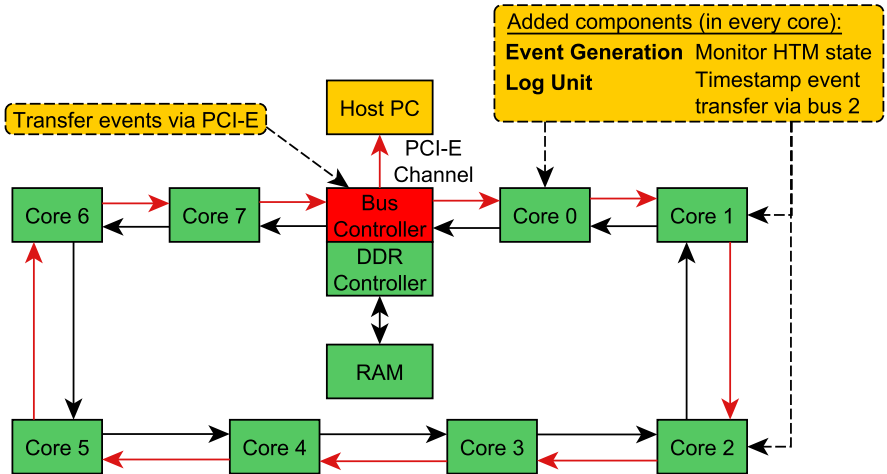
Design

Event Generation, Log Unit, Bus Controller



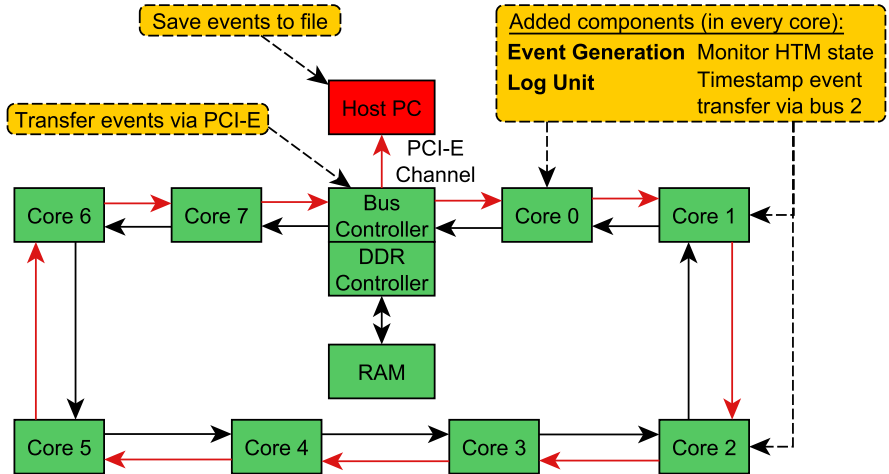
Design

Event Generation, Log Unit, Bus Controller



Design

Event Generation, Log Unit, Bus Controller



Design

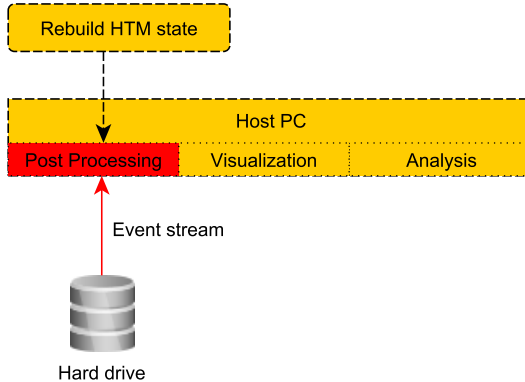
Post Processing, Visualization, Analysis



Hard drive

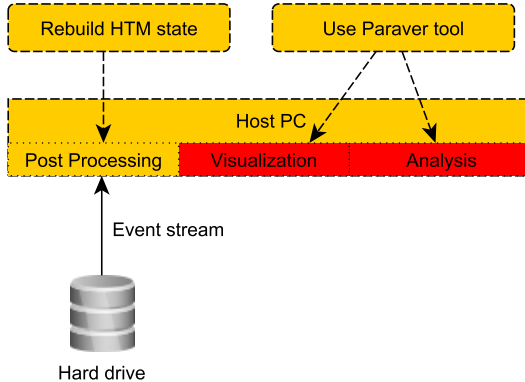
Design

Post Processing, Visualization, Analysis



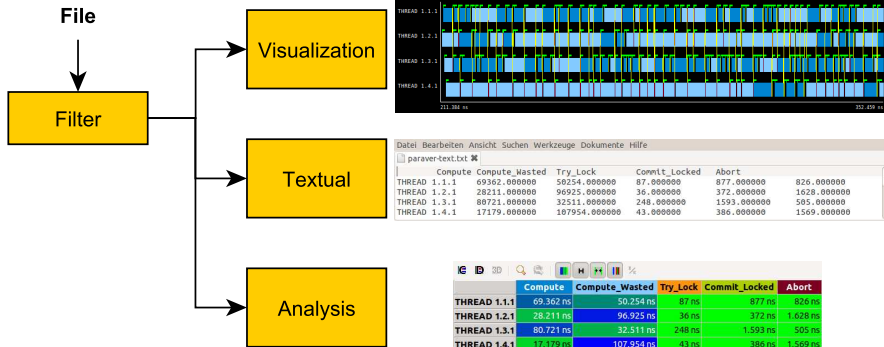
Design

Post Processing, Visualization, Analysis




Results

Paraver workflow



Results

Example: 4 conflicting threads



Idle



Compute



Compute Wasted



Try Lock



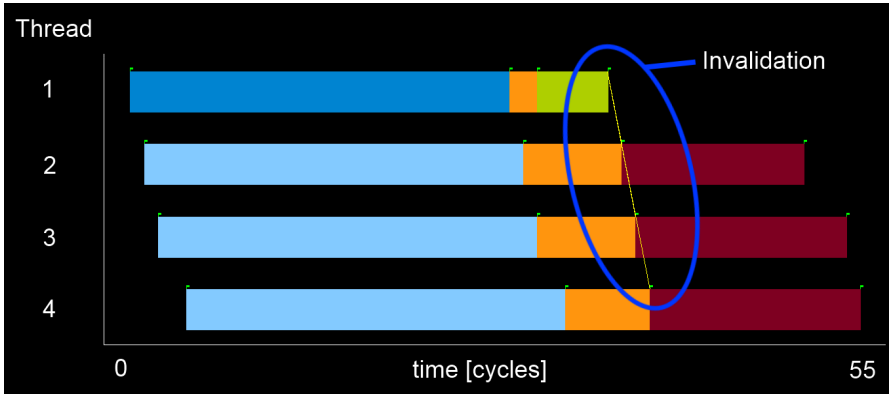
Commit



Abort

Results

Example: 4 conflicting threads



- ✓ Application / HTM runtime behavior is trackable with no probe effects
- ✓ Visualization capabilities lead to in-depth understanding of application / HTM runtime behavior
- Metrics provided via offline analysis, e.g.
 - Time spent in Committed & Aborted Transactions
 - Contention / Commit & Abort Rate
 - Contention between specific threads
 - HTM System Overhead

Summary at a glance

The TMbox system now supports:

- ✓ Identification of detailed runtime characteristics of an application
- ✓ Easy detection of application bottlenecks
- ✓ Getting hints to optimize application concerning both performance and scalability

Ongoing work

Combine monitoring with STM runtime environment:
Allows analysis of HybridTM systems

See further work in full paper “A low-overhead profiling and visualization framework for Hybrid Transactional Memory” (to appear in FCCM 2012)

Acknowledgment



This work was supported by a Short Term Scientific Mission (STSM) grant from the European Cooperation in Science and Technology (COST) Action IC1001 (EuroTM) during June and July 2011.

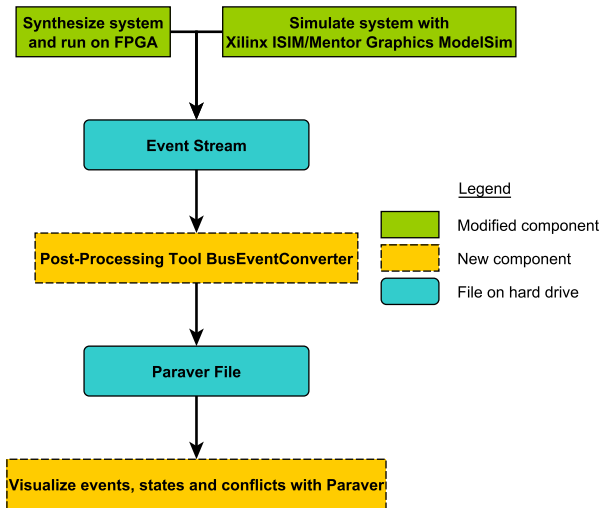
Enhancing an HTM System with Monitoring, Visualization and Analysis Capabilities

**Philipp Kirchhofer¹, Martin Schindewolf¹, Nehir Sonmez², Oriol Arcas²,
Osman S. Unsal², Adrián Cristal², Wolfgang Karl¹**

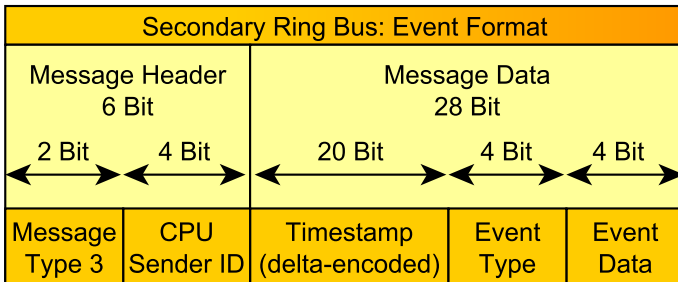
¹ Karlsruhe Institute of Technology (KIT) ² Barcelona Supercomputing Center (BSC)

Chair for Computer Architecture and Parallel Processing, Institute of Computer Engineering, Karlsruhe Institute of Technology


Additional slides



Event Diagram



Event diagram



TMbox FPGA Usage (with monitoring infrastructure) - Increase per Core