




The GACOP Research Group

José M. García
Computer Engineering Department
University of Murcia (Spain)


jmgarcia@dittec.um.es



Hipeac-2 Kick-off meeting
Gothenburg, January 30-31, 2008

Context: Virtualization

- Virtualization techniques are becoming increasingly important for fully utilizing future many-core systems.
- Recent commercial architecture advances, such as Intel's VT or AMD Pacifica technologies.
- CMPs provide excellent new opportunities to expand server and workload consolidation for VM.
- It will be necessary to design a memory hierarchy to allow greater use of space sharing.




Challenges and our current work

- The memory system should
 - Maximize shared memory accesses serviced within a VM
 - Minimize interference among separate VMs
 - Facilitate dynamic reassignment of cores, caches, and memory to VM
 - Support sharing among VMs
- Our proposal
 - It is based on Virtual Hierarchies (Hill's paper on last ISCA)
 - We'd like to apply our Direct Coherence protocol to improve both performance and implementation ease

Our proposal: Direct Coherence

- Goals:
 - To avoid the indirection of directory-based protocols without relying in broadcast or any other brute-force method.
- Direct Coherence
 - The role of ensuring totally ordered accesses for every memory block is assigned to the owner node, instead of the home node.
 - Indirection is avoided by directly sending the requests to the owner node, which can provide the up-to-date copy of the block.
 - Improvements in execution time of 8% (on average) over a directory protocol.
 - Reductions in network traffic of 42% (on average) compared to Token.

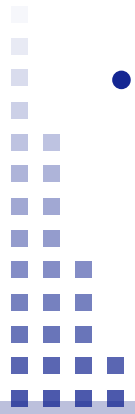


Our expertise: Topics

- Computer Architecture (Branch Prediction and Value Prediction, Multithreaded Architectures)
- Low power techniques for high performance processors and embedded processors
 - Reducing the energy consumption of the front-end for single-chip architectures
- Shared-memory multiprocessors (cache-coherent non-uniform memory access architectures or cc-NUMAs)
 - A novel directory architecture —multilevel directories—
 - On-chip directory (including the first-level directory, along with a shared-data cache, into each one of the processor chips)
 - Removing the directory latency from the critical path for both cache-to-cache transfer misses and upgrade misses

Our expertise: Tools used

- Detailed execution-driven simulators
 - SimpleScalar
 - Wattch
 - RSIM
 - Wisconsin GEMS

- 
- Functional simulators
 - Simics
 - CACTI
 - HotLeakage, HotSpot