

Adaptive Compilation Cluster

Paris 27/Nov/2008

S. Bartolini – Università di Siena, Italy - bartolini@dii.unisi.it

Topics I am currently working on:

- ▶ Feedback-directed post-link program restructuring to reduce cache power consumption
 - ▶ Tagless accessible cache (with University of Edinburgh, University of Gent)
 - ▶ In conjunction with other Hw mechanisms (e.g. drowsy/decay mechanisms)
- ▶ Feedback-directed compiler optimizations for L2/L3 NUCA Cache management (with University of Pisa)
 - ▶ Performance and power in multicores and parallel applications
- ▶ Runtime/Compiler-assisted cache/scatchpad mapping/tuning for power management: modulate resource usage according to application needs
 - ▶ (One of my students visited ARM on topics in this line)
- ▶ Profiling/tracing of multithreaded applications as to allow insight in their runtime behavior and enable new optimizations and tuning
Integrated tracing of O.S. activity: thread scheduling and interaction with Hw res. (some collaboration with UPC)
 - ▶ E.g. x264 h.264 encoder on real multicore machines (Intel Core 2 Duo, ...)
- ▶ Inclusion of cryptographic accelerators inside embedded processors (U. Siena)
 - ▶ Design space exploration on the configuration, proc/accel interfacing, accel ISA
 - ▶ Gcc for compiling applications supporting the accelerator instructions
- ▶ Elliptic-curve cryptography (Public-key cryptosystem suitable for mobile devices)

Topics I am interested in ...

... are not limited to 😊

- ▶ Synergistic *software*(compiler/linker/loader/runtime)-*hardware* proposals aiming to optimise the significant end-user metrics in specific application domains (e.g. performance/power/predictability)
 - ▶ Parallel programming: models, auto-parallelization, optimizations, sw/hw TM
 - ▶ Multicore (homogeneous/heterogeneous) architectures, domain-specific accelerators
 - ▶ Articulated memory hierarchies (caches, scratchpad), NoCs
 - ▶ Reconfigurability
-
- ▶

Sample topics for stimulating our discussion

- ▶ Power implications (and opportunities) of parallelization approaches in multicores
 - ▶ At all levels: coding, compiling (auto-parallelization), runtime
 - ▶ Tuning to specific applications (e.g. embedded)
- ▶ Applications that are able to automatically issue *tuning* commands to the architectural (cores, accel., memories, NoC,...) and non-arch (O.S., runtime, ...) modules as to adapt to the runtime conditions of the system within the overall objectives (performance/power tradeoffs)
 - ▶ Feedback-directed optimizations
 - ▶ Parametric models of the architecture elements and prediction of tuning effects
 - ▶ Global application and architecture view



How to interact ?

- ▶ I am available for collaboration on these, related and complimentary topics:
 - ▶ Exchange ideas and explore possible proposals for applying to EU FP7 call
 - ▶ Get together and work on a topic (e.g. joint paper perspective)



Thanks !

Sandro Bartolini
Università di Siena

bartolini@dii.unisi.it

