

Hardware Group

NSF Workshop on the Science of
Power Management

Hardware (Synopsis: Fund Us)

- Tradeoffs: Power vs performance vs reliability (Also: thermals & variability)
- Need models to reason about designs at many levels:
 - Levels: Circuits, chips, systems, data centers, power distribution, ...
 - Issues: Performance, power, reliability, variability, ...
 - Composition of models: hierarchy, interfaces, co-simulation of models at different levels of detail, information flow across domains, ...
 - Unsolved bits: unification of “point solutions”, multi-scale models/simulators
 - Needs constant attention: Verification, validation, scalability
- Need models to teach programmers to reason about energy efficiency
 - Something akin to $O(*)$ notation for energy
- Hierarchical/multi-level control systems for power/thermal management
 - Separate control systems/agents managing at diff. levels (circuit, chip, system, ...)
 - What function should be at each level? What should be in HW vs FW vs VM vs ...
 - How do these control systems interact? Autonomic/self-coordinating?
 - What levels should individually be “power proportional”?
 - Need test bench, workloads, repeatability for analysis

Hardware

- What is the role of heterogeneity?
 - How much heterogeneity (if any)? What range of power-perf-... tradeoffs? Same ISA? Von Neumann vs stream vs vector vs special-purpose accelerators?
 - Single-threaded performance vs throughput performance
 - What is heterogeneous? Wires? Circuits? Functional units? Processors?
 - How do you match form of heterogeneity to set of applications you want to run?
 - How do heterogeneous components interact?
- Emerging technologies: (How do we integrate/exploit it?)
 - Memory: Embedded DRAM, Flash, PCM, sleep-modes in DDR3/4, ...
 - Interconnects: optical, copper, on-chip/off-chip
 - Exploit technologies to broaden range of power/performance tradeoffs:
 - Dynamic (active) range, sleep options, ...
 - 3D integration
 - Multiple voltage and power domains (and interfaces between them)
 - Power distribution network
 - Multiple on-chip VRMs (for dynamic scaling within domains)

Source Slides

Hardware (H1)

- Power-reliability-performance tradeoffs (thermals?)
 - Reliability issues: Aging, soft errors,
 - How much extra energy do you need to get the next higher level of reliability?
 - Science issues:
 - What should we optimize? (Power, power-performance, power-perf², ...)?
 - Once we have defined our constraints and goals, solve the optimization problem...
- Need hierarchical “controls” (interfaces) that control tradeoffs between power, performance, and reliability (and maybe more)
 - Chip, board, rack, system
 - What should run in hardware vs firmware vs hypervisor vs OS vs application vs independent (autonomous) software systems
 - At what levels should systems be energy-proportional?
 - Science issues:
 - How do we manage these autonomous systems? Control theory (“self-coordinating control systems with limited signaling”)
 - Define a hierarchical model of how these systems are managed and reason about what controls are needed
 - What are the tradeoffs of doing the management at a given level? (Power/perf/reliab.)

Hardware (H1)

- Need to develop principles for programmers (akin to $O(*)$ for performance) that guides people to write less energy-inefficient applications
 - Is there a simple equation/model that we can use to reason?
- What is the role of heterogeneity?
 - How do we exploit the heterogeneity? Same/different ISA? Power ranges?
 - How broad a range of heterogeneity (if any) makes sense?
 - Same or different ISA? If same ISA, maybe exploit the AI ideas where they write multiple algorithms and select the appropriate one at run-time
 - Von Neumann vs stream vs vector vs special-purpose accelerators?
 - Do the “other” issues (like verification, coherence, etc.) make this infeasible?
 - Single-threaded performance vs throughput performance
 - Heterogeneous... Wires? Circuits? Functional units? Processors?
- What about memory? (Elephant in the room)
 - Different types of RAM (or circuits) with different power-perf-reliability tradeoffs
 - Differently clocked DRAM, Flash, phase change memory (PCM)
 - Gating / compression memory → who manages it?
- What kind of sleep modes (c states)? Who manages and how?
 - Why can't manufacturers agree to common standard (“standard” ACPI)?

Hardware (H1)

- Can we even power all of the processor/transistors at the same time?
 - “Nominal” does not even have a clear meaning any more
- What can we do to support lower supply voltages?
 - Difficulties: SRAMs, reliability, ...
- What will the form factor will be used to interface end-users?
 - Desktops vs laptops vs netbooks vs PDAs vs cell phones
 - Can we dramatically reduce the power burned by these end systems?
- Can “science” play a role in reducing our large safety/design margins?
- Technologies:
 - 3D integration
 - On-/off-chip optical interconnects
 - Storage-class memories
 - Non-CMOS technologies
- How do we undo the damage of the “every problem in CS can be solved with another layer of indirection” phenomenon?

Hardware (H1)

- What roles can/should “offload processors” play?
 - Wake-on-LAN (offload ARP, SSH keepalive, ...)
 - Remote Disk/RAM access (RDMA, shmem, ...)
- Need support for building prototype chips/systems for studying low power issues
 - The amount of hardware design research in the US has dropped dramatically
- Lack of sufficient benchmarks, models, large-scale/power simulators
 - What role should abstract models play vs simulators?
- Science questions:
 - Can we answer questions regarding what degree of heterogeneity is worthwhile, what (rough) potential benefits would we achieve, etc.?

Hardware (H2)

- Circuits that support:
 - Different energy-delay tradeoffs (heterogeneity at different levels)
 - Sleep states (with model of leakage/sleep power vs latency)
- Multi-tiered (hierarchical) power and thermal management
 - High-level PTM accepts user/admin input, does coarse-grained optimization
 - Low-level PTM accepts constraints from above, manages details of
 - Multi-variate optimization: security, fault tolerance, power, ...
 - Game theory (?), control theory
 - Benchmarks, model synthesis, extracting models on testbed to characterize them (in particular, need workloads that drive full hierarchy)
 - Component-level models and composite model that account for application and hardware characteristics
 - Virtualization (already being pushed hard at enterprise level)
- Heterogeneity
 - What kind of heterogeneity? How do you match form of heterogeneity to set of applications you want to run? How do heterogeneous components interact?

Hardware (H2)

- Handling manufacturing-induced forms of “uncertainty”
 - Take existing techniques that are used at the circuit/technology level and apply them at higher levels
- On-chip power distribution mechanisms
 - Hard: On-chip regulators, coupling caps, interactions between voltage domains
 - Could use models to reason about value and complexity of adding any given level of heterogeneous voltages and voltage domains
 - Models for reasoning about power delivery
- Need fast, chip-level profiling tools to be able to predict thermal impact of making various run-time decisions
 - Doing it accurately is very computationally-intensive (e.g., CFDs)
 - Doing it super-fast can be very inaccurate (heuristics)
 - Can we develop a “happy medium”?
 - Models for reasoning about thermal issues
- Energy-efficient interconnects (both on-chip and off-chip)
- Integrated model-of-models to allow reasoning about entire set of problems:
 - power, power delivery, thermals, heterogeneity, hierarchy, ...