

ECE 571 – Advanced Microprocessor-Based Design Lecture 15

Vince Weaver

`http://web.eece.maine.edu/~vweaver`

`vincent.weaver@maine.edu`

30 March 2017

Announcements

- HW#7 Due
- Project ideas due



Missed topics from Last time

- Why does kernel live in address space? speed (20%) and can access all pointers? Also could be security issue
- TLB shutdown. If a page remapping happens on one processor, have to change the mapping across all processors in system
- TLB knows about permission bits (r/w/x). How does cache know if it's allowed.



Haswell Virtual Memory

- L1 (4-way associative)
 - 64 4kB
 - 32 2MB
 - 4 1GB
- L2 (1024 entry 8-way associative, combined 4kB and 2M)
- DCache – 32kB/8-way so VIPT possible



Cortex A9 MMU

- Virtual Memory System Architecture version 7 (VMSAv7)
- page table entries that support 4KB, 64KB, 1MB, and 16MB
- global and address space ID (no more TLB flush on context switch)
- instruction micro-TLB (32 or 64 fully associative)



- data micro-TLB (32 fully associative)
- Unified main TLB, 2-way, 2x64 (128 total) on pandaboard
- 4 lockable entries (why want to do that?)
- Supports hardware page table walks



Cortex A9 MMU

- Virtual Memory System Architecture version 7 (VMSAv7)
- Addresses can be 40bits virt / 32 physical
- First check FCSE – linear translation of bottom 32MB to arbitrary block in physical memory (optional with VMSAv7)




Cortex A9 TLB

- micro-TLB. 1 cycle access. needs to be flushed if ASID changes
- fully-associative lockable 4 elements plus 2-way larger. varying cycles access



Cortex A9 TLB Measurement




CPU Power and Energy

- Became a trendy thing to research in 1999-2002 timeframe.
- Before that usually concern was with performance.
- These days energy results are often reported as a core part of any architectural proposal, not as a separate issue.
- The results discussed here are academic and may or may not be implemented in actual chips.



AMD Bulldozer Die Shot



Note which structures are big, using static power.



CPU Power Breakdown

From Fan, Tang, Huan, Gao (ISLPED'05), Chinese Godson MIPS CPU

They gave numbers, but unclear of workload, if static or dynamic, etc.

- Cache 36%
- TLB 13%
- FALU 10%
- ROQueue 7%
- FMUL 6%



- Float reg 5%
- Gen reg 5%
- MUL 2%
- MCUControl 2%
- ALU 1%
- Other 13%



Thermal Concerns Too

Power density exceed hot plate, approaching rocket nozzle

TODO: Find the Intel cite for this statement.



Methodologies Used in These Papers

It varies, but many of these are from simulations (sometimes validated). Anything from SPICE to “cycle-accurate” simulators.



Clock Generation

- Driving high-frequency load against capacitance, trying to keep whole chip in sync.
- Extreme Case: Alpha 21264 H-tree, 32% of power?
- Half-frequency clocks (on both edge, so clock run half as fast) (Mudge 2001)
- Asynchronous
- Locally Asynchronous (Divide to multiple clock domains)



DVFS and other CPU Power/Energy Saving Methods

- A lot of related work
- Will focus on actual implementations rather than academic papers this time



DVFS

- Voltage planes – on CMP might share voltage planes so have to scale multiple processors at a time
- DC to DC converter, programmable.
- Phase-Locked Loops. Orders of ms to change. Multiplier of some crystal frequency.
- Senger et al ISCAS 2006 lists some alternatives. Two phase locked loops? High frequency loop and have programmable divider?



- Often takes time, on order of milliseconds, to switch frequency. Switching voltage can be done with less hassle.



Adaptive Body Biasing

- Related to but not always considered part of DVFS
- Control voltage applied to body
- Change the threshold voltage
- Reduces leakage but slows performance



DVFS and other CPU Power/Energy Saving Methods

- A lot of related work
- Will focus on actual implementations rather than academic papers this time



DVFS

- Voltage planes – on CMP might share voltage planes so have to scale multiple processors at a time
- DC to DC converter, programmable.
- Phase-Locked Loops. Orders of ms to change. Multiplier of some crystal frequency.
- Senger et al ISCAS 2006 lists some alternatives. Two phase locked loops? High frequency loop and have programmable divider?



- Often takes time, on order of milliseconds, to switch frequency. Switching voltage can be done with less hassle.




Adaptive Body Biasing

- Related to but not always considered part of DVFS
- Control voltage applied to body
- Change the threshold voltage
- Reduces leakage but slows performance



Cache Power and Energy

Large area, low-hanging fruit



Decay Caches

- Kaxiras, Ho, Martinosi (ISCA 2001)
- Turn off cache lines not being used to reduce leakage
- DRAM cache with no refresh
- Decayed values can be re-fetched from memory.
Tradeoff.



Drowsy Caches

- Flautner, Kim, Martin, Blaauw, Mudge. ISCA 2002.
- Move cold cache lines into “drowsy” mode.
Lower power enough to hold state, not enough to lose contents. Reduce leakage. Better than decay as not lose data.
- Note: in Intel Volume 3b 17.17.5.2 it mentions certain C states might power down or otherwise turn off parts of cache.



Adaptive Caches

- Albonesi (Micro 1999). Manually turn off ways in cache with an instruction.
- Size the caches



Cache Compression

- Dynamic zero compression for cache energy reduction (L Villa, M Zhang, K Asanović. Micro 2001).
- Cache Compression (“sign compression” – top bits)
Energy savings 20% (simulated) (Kim, Austin, Mudge
WMPI 2002)



Banking and Filtering

- Filter cache, banking (only have half of cache active) (Mudge 2001)
- Slowing Down Cache Hits, Banked Data Cache. (Huang, Renau, Yoo, and Torrellas. Micro 2000.)
- Vertical Banking, Horizontal Banking (Su and Despain, ISLPED 1995).



Code Scheduling

- Can Schedule code for lower power.
- Better cache rates lower power. performance/power can go hand in hand. (Kandemir, Vijaykrishnan, Irwin)



Branch Predictors

- Parikh, Skadron, Zhang, Barcella, Stan
- 4 concerns:
 1. Accuracy. Not affect power, but performance
 2. Configuration (may affect power)
 3. Number of lookups
 4. Number of updates
- Tradeoff power vs time.



- brpred can be size of small cache, 10% of power
- Can use banking to mitigate



Branch Predictors

- can watch icache, not activate predictor if nobranches
- Pipeline gating, keep track of each predicted branch confidence. If confidence hits certain threshold, stop speculating. Show this may or may not be good.
- Integer code, large predictors good
- FP, tight loops, predictors not as important.



Branch Predictor Evaluation

- (Strasser, 1999). Simulation, small branch predictor can help energy.
- (Co, Weikle, Skadron) Formula for break even point. Leakage matters, what brpred hides is stall cycles.
- SEPAS: A Highly Accurate Energy-Efficient Branch Predictor (Baniasadi, Moshovos. ISLPED 2004).
Once a branch prediction reaches steady state (unlikely to change) stop accessing/updating predictor, saving



energy.

- Low Power/Area Branch Prediction Using Complementary Branch Predictors (Sendag, Yi, Chuang, Lija. IPDPS 2008)

Complementary Branch Predictor to handle the tough cases.



Prefetching


- Prefetching does not get looked at as closely. Various studies show it can be a win energy wise, but it is a close thing.
- (Guo, Chheda, Koren, Krishna, Moritz. PACS'04)
HW Prefetch increase power 30%; have compiler help augment with hints, filters.
- (Tang, Liu, Gu, Liu, Gaudiot. Computer Architecture Letters, 2011).



Mixed results.



TLB Energy



TLB Optimization – Assume in Same Page

- Optimizing instruction TLB energy using software and hardware techniques (Kadayif, Sivasubramaniam, Kandemir, Kandiraju, Chen. TODAES 2005).
Don't access TLB if not necessary. Compare to last access (assume stay in same page) Circuit improvements
- (Kadayif, Sivasubramaniam, Kandemir, Kandiraju, Chen. Micro 2002)
Generating Physical Addresses Directly for Saving Instruction TLB Energy Cache page value.



TLB Optimization – Use Virtual Caches

- (Ekman and Stenström, ISLPED 2002) Use virt address cache. Less TLB energy, more snoop energy. TLB keeps track of shared pages.



TLB Optimization – Reconfiguring

- (Basu, Hill, Swift. ISCA 2012) Reducing Memory Reference Energy with Opportunistic Virtual Caching
Have the OS select if memory region physical or virtual cached.
- (Delaluz, Kandemir, Sivasubramaniam, Irwin, Vijaykrishnan. ICCD 2013) Reducing dTLB Energy Through Dynamic Resizing.
Size TLB as needed, shutting off banks. Easier if fully-associative.



TLB Optimization – Memory Placement

- (Jeyapaul, Marathe, Shrivastava, VLSI'09) Try to keep as much in one page as possible via compiler.
- Energy Efficient D-TLB and Data Cache using Semantic-Aware Multilateral Partitioning (Lee, Ballapuram. ISLPED'03) Split memory regions by region (text/data/heap). Better TLB performance, better energy.



Bus Protocols

- Bus Protocols
- Cache-Coherence Protocols



Busses

- Grey Code, only one bit change when incrementing.
Lower energy on busses? (Su and Despain, ISLPED 1995).

