



A Power Analyzer for Pocket Computers



■ Deliverable: PowerAnalyzer

- a power aware cycle-level simulator

■ Personnel

- Todd Austin

- Advanced Computer Architecture Lab, The University of Michigan, Ann Arbor

- Dirk Grunwald

- Department of Computer Science, The University of Colorado, Boulder

- Trevor Mudge

- Advanced Computer Architecture Lab, The University of Michigan, Ann Arbor and ARM Computers, Cambridge England

■ Start date: June 1, 2000

■ Duration: 2 years





- **Leverage existing work and research**
 - SimpleScalar cycle-level simulator (NSF CADRE - Todd)
 - energy efficient OS (NSF - Dirk)
 - low power embedded processors (NSF- Trev)
- **Open development model**
- **Modular design**
 - plug-ins for various levels of abstraction and modeling techniques
- **Supported by the SimpleScalar distribution/maintenance model**





PI's Backgrounds



■ Until recently

- High performance computer architecture
- Microarchitectures
- Processor design
- Operating systems
- Compiler/architecture trade-offs
- Performance evaluation and simulation

■ Recently

- Workshops/Tutorials on power aware computing (ISCA98/MICRO99)
- Embedded computers
 - **Design and programming**
 - **8-bit low power**
 - **Code compression**
- Pocket computers
 - **ITSY — Compaq WRL**
 - **System software and JVMs**





Goal for PowerAnalyzer



- **Make power a 1st class design consideration like performance — this implies we need:**

- ☐ early power estimates during performance studies
- ☐ combine into cycle-level simulators

- **Cycle-level simulator**

- ☐ simulates activity or events in the microarchitecture
 - **pipelines; caches; memories, decoders, function units**
- ☐ more detailed than an ISA simulator
- ☐ less detailed than a register-transfer simulator





- Power/energy are important for portables because of battery life and weight (obvious)

- Equally important for most other classes of computing too

- Compaq's Alpha

- 100 A @ 150 W

- Power density in $\text{W}/\text{cm}^2 = 33$

- Hot plate = 10
 - Nuclear reactor = 200

- Server farm: Heavy duty factory

- 25,000 sq. ft.

- ~8,000 servers

- ~2,000,000 Watts

	Power (Watts)	Freq. (MHz.)	Die Size (mm ²)	Vdd
--	------------------	-----------------	--------------------------------	-----

Alpha 21064	30	200	234	3.3
----------------	----	-----	-----	-----

Alpha 21164	50	300	299	3.3
----------------	----	-----	-----	-----

Alpha 21264	90	575	314	2.2
----------------	----	-----	-----	-----

Alpha 21364	>100	>1000	340	1.5
----------------	------	-------	-----	-----





Deliverable: PowerAnalyzer



■ Develop a practical power evaluation tool

- suitable for calculating power consumption for computer systems
- calibrated against physical targets

■ Initial target

- StrongARM hand held with full linux OS
- ITSY - iPAQ 3600 commercial version
- pocket computers where computing and communication place strong demands on the portable power supply

■ Capabilities

- dynamic and static power consumption profiles
- power/performance trade-offs
- complete systems — processors, memories and peripherals
- preserve power ordering





Power Analysis: What Questions?



- **Average (static) power consumption**
 - battery life
- **Dynamic power: di/dt**
 - inductance effects
- **Peak power**
 - limit violations
- **Voltage scaling/clustering**
- **Thermal guestimates**
- **How do these quantities change as we change the microarchitecture?**





At What Abstraction do we Perform Power Analysis?



- **Layout: accuracy ~5%**
 - extract parameters => SPICE
- **Logic-level: accuracy ~10%**
 - combine: toggle counts, fan-out, leakage, dc, short circuit
- **RTL: accuracy ~ 20%**
 - based on activity counts (Watt Watcher by Sente)
- **Cycle level: accuracy ~30% or less**
 - useful for relative measures/flag potential peaks (should preserve ordering at least)





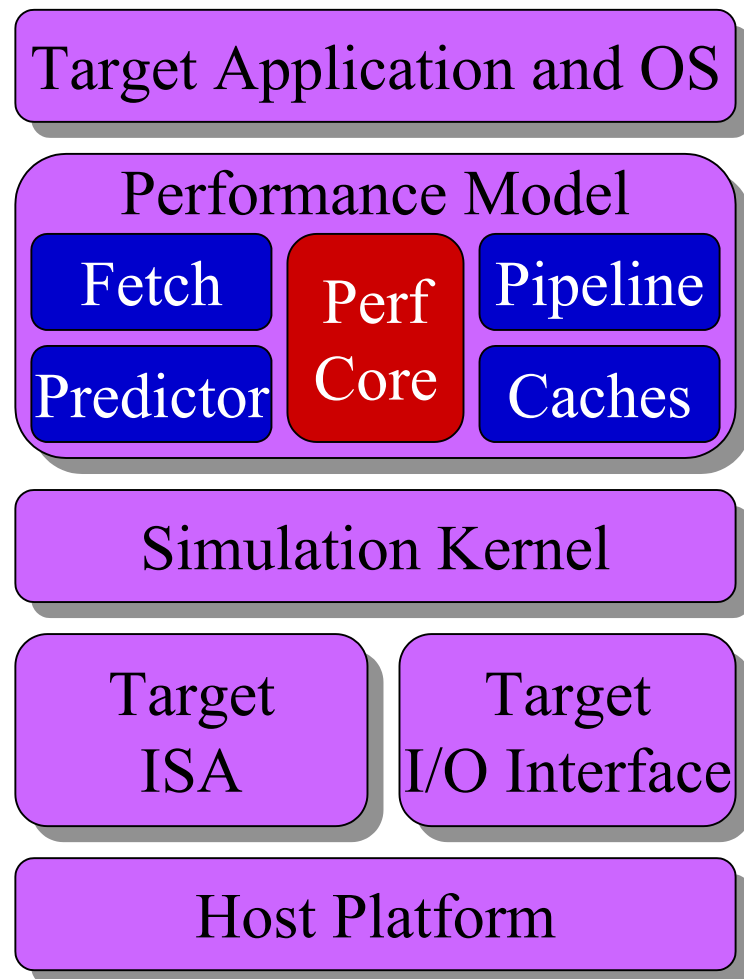
Approach



- **Base on SimpleScalar an existing widely used cycle-simulator**
- **Computer system design and analysis infrastructure**
 - processor and device models
 - support for popular ISAs and I/O interfaces
 - simulators run on most modern platforms
- **Created by the SimpleScalar development team**
 - UM, UW-Madison, UT-Austin
 - entering eighth year of development
 - in use at many universities and companies
 - in '99, nearly 3 out of 10 papers in architecture conferences used SimpleScalar
- **Freely available with source and documentation**
 - www.SimpleScalar.org



- **Target application and operating system run on simulator**
- **Performance model tracks time**
 - performance core implements machine
 - many standard modules to speed development
- **Simulation kernel provides discreet event simulation services**
- **Target ISA support**
 - PISA, Alpha AXP, PPC
- **Target I/O support**
 - syscalls, devices, I/O traces





What “Power” Metric



- **Total energy**
 - **correlated with battery life**
 - BUT battery life also a function of usage
 - usually reduced by high power use
- **Total energy with a deadline**
- **High performance designers often interested in peak power limits**
- **Another important limit: di/dt noise**
- **MIPS/W**
 - **work per W — energy used per average instruction**





Adding Power



- **The key challenge — greatest potential for innovation**
- **Current (?) approaches use frequency counts of “interesting” nodes in the microarchitecture**
 - obtained from running benchmarks through simulator
- **Is this sufficient?**
- **Do we need to run simulations to completion?**
 - can we sample?
- **Balance speed against accuracy**
 - e.g. can we omit data dependency considerations
- **Guideline: preserve power ordering**
 - if $\text{power}_1 > \text{power}_2$ then $\text{sim_power}_1 > \text{sim_power}_2$
- **Can we deduce useful information by comparing the frequency count vectors**





Adding Power: Initial Idea



- **Subdivide cycle simulator into blocks**
 - fetch unit, tlb, cache tags, etc.
 - composed of library modules
- **Track events in these blocks using counters in SimpleScalar**
 - obtain a block level frequency measure
 - important to measure activity that doesn't help forward progress of the computation
- **Blocks are characterized by:**
 - active and inactive power density x equivalent gate count
- **Dynamic power obtained from cycle-to-cycle changes**
- **Proof of concept:**
 - **G. Z. Cai and C. H. Lim of Intel's Mobile CPU group**
 - Architectural Level Power/Performance Optimization and Dynamic Power Estimation. In Proc. Cool Chips Tutorial, MICRO32, Haifa Israel, pp. 90-113.





- **SPICE simulations of existing UM prototype cores and components**
 - ARM; PPC; function units; pipelines
 - library modules
- **Measurements of running ITSY Pocket Computer**
 - voltage scaling
- **StrongARM based**
 - extend SimpleScalars's processors models to include ARM ISA





Schedule



		2000			2001						
		July/Aug	Sept/Oct	Nov/Dec	Jan/Feb	Mar/Apr	May/June	July/Aug	Sept/Oct	Nov/Dec	Jan/Feb
Michigan											
	Functioning Simulator								Voltage Scheduling Policy		
			O/S Simulator								
	Functioning Simulator						Pipeline Gating				
			SA-1100 Simulator						Memory S		
Colorado	Component Design								Initial Whitebox Calibration		
					MARM Simulator			Advanced			
	Component Design					Blackbox Calibration					
			SA-1100 Simulator								Cluster
	Voltage Scheduling in Linux								Voltage Scheduling		
					Application Event Interface						





- **Memory: 1\$; 2\$; 3\$... ; PM ; disk**
- **Growing component of chip size**
 - Alpha '264: 64K/64K 1/2 the die area
 - HP PA-8500: 512K/1M
 - M.Core - most of the die
- **Reduce activity in memories**





Reducing Memory Activity



■ Shield cache/filter cache

□ intercepts addresses and accesses to \$

- 50% hit rate ok, so it can be small (128 lines)

■ Banked cache

□ require references to remain in a bank

□ instruction references exhibit this

- sojourn in 128 line banks was significant
- not so good for data caches
- tolerate an extra cycle to switch banks





Reducing Memory Activity



- **Make 2\$ victim cache – no duplication**
- **Associativity is not good**
 - too much activity on comparators
- **Amulet (asynchronous ARM) stood this on it's head**
 - use the “ways” as banks and keep active the MRU bank/way (requires prediction)
- **MRU can reduce effects of associativity**
 - associative caches can be smaller





Reducing Memory Activity



■ DRAM on the chip

- IRAM project
- Embedded processors – Mitsubishi 32R/D
 - 32-bit risc + 2 MB DRAM 100MHz 1997

■ Saves power

- 1T cell vs 4-6T cell
- 10-1 density — IBM SA-27E asic process
 - Less chip crossings

■ NeoMagic's graphics controller 500-750 mW

- 25% of a multi-chip alternative





Trends Against Low Power



■ Speculation for high performance

- ☐ branch prediction
- ☐ trace caches
- ☐ prefetching
- ☐ runahead
- ☐ etc
- ☐ CPUs not important?

■ Interpretation more popular again

- ☐ JVMs
- ☐ can compilers help us with JITs?
- ☐ can hardware assists help?





■ Frequency scaling — power is linear in frequency

- often not effective alone because correspondingly less work can be done
- useful when processor is idle (next slide)
 - frequency = 0!
- useful when process is rate determined
 - mpeg requires 30 frames / second and no more

■ Voltage scaling — power is quadratic in voltage

- accompanied by reduction in frequency
- but reduction in power is quadratic

■ Scheduled by OS (clock gating is an alternative)

- limited by PLL settling time



■ Pentium III SpeedStep

- 650 MHz consumes 14.4 watts of power at 1.6 volts
- 500 MHz consumes 7.9 watts of power at 1.35 volts
- wall plug or mobile

■ StrongARM SA-2

- 450 mW @ 600 MHz 1.3V
- 40 mW @ 150 MHz 0.75V
 - Dhrystone 2.1



Idle Time in Desktop Applications



Category	Application	Idle
Automated	Netscape - sys98	17.80%
	Lotus Office - wins99	48.60%
	Microsoft Office - wins 99	17.80%
	Visual C++ - wins99	16.00%
	Photoshop - sys98	24.20%
	xmcode - sys98	15.90%
	Microstation - wins99	17.70%
	Paradox - sys98	19.20%
Interactive / Multimedia	Windows Media Player - AVI	73.50%
	Windows Media Player - MP3	92.80%
	QuickTime - Video	44.40%
	QuickTime - MP3	92.80%
	HotJava	78.00%
	Internet Explorer	90.50%
	Cloudscape	29.70%
	Image/J	42.20%
	PowerTranslator	29.40%

*Done with Kris Flautner UM and Rich Uhlig Intel Microprocessor Research Lab
using a 4-processors Pentium III system.





■ Voltage scheduling is being supported in newer processors

- recall that $1/2$ voltage reduces the power to 25%, but f is reduced only by 50%
- SA-2 will support dynamic voltage scaling (earlier)

■ Turn off the clock

- requires static design
- low leakage too





Summary



- **Leverage existing work and research**
 - SimpleScalar cycle-level simulator (NSF CADRE - Todd)
 - energy efficient OS (NSF - Dirk)
 - low power embedded processors (NSF- Trev)
- **Open development model**
- **Modular design**
 - plug-ins for various levels of abstraction and modeling techniques
- **Supported by the SimpleScalar distribution/maintenance model**

