Processor-Memory Power Shifting

Workshop on Energy-Efficient Design
June 9, 2012

Heather Hanson, Wes Felter, Wei Huang, Charles Lefurgy,
Karthick Rajamani, Freeman Rawson, Guillermo Silva

Outline

Power profiles

- Architectural evolution
- Effects on power management

Power-shifting study

- Policies
- Results
- Discussion

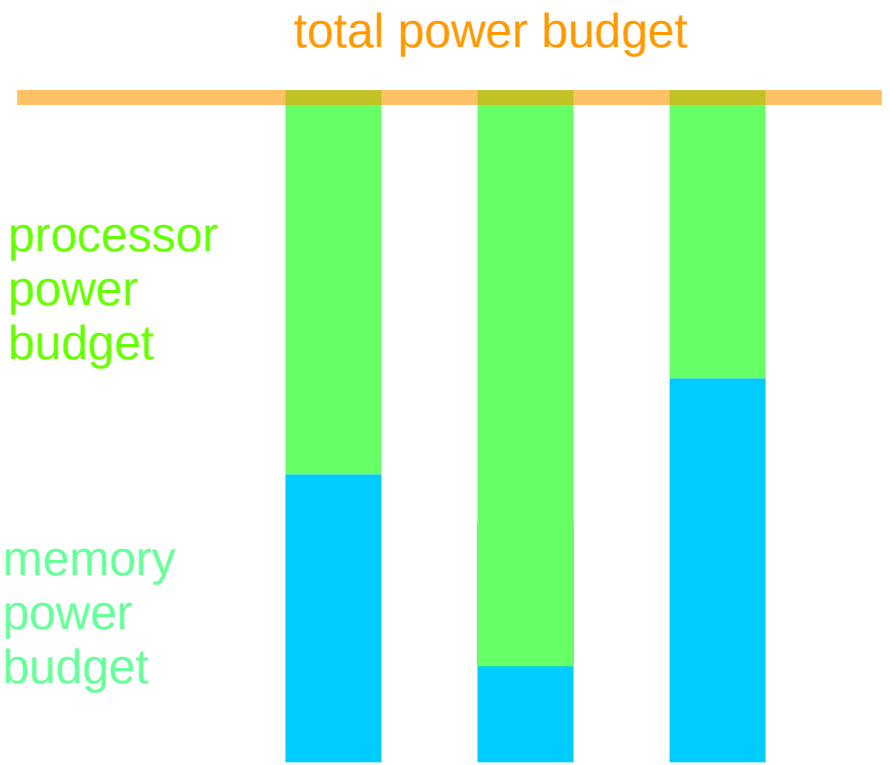

Before.....

inverse relationship between processor and memory power for single-core, single-threaded, stalling pipeline

Power shifting opportunity:

- scale back memory power budget when processor is active, and vice versa
- minimal performance effects because other component doesn't need the power right now anyway

iso-power line
50W total

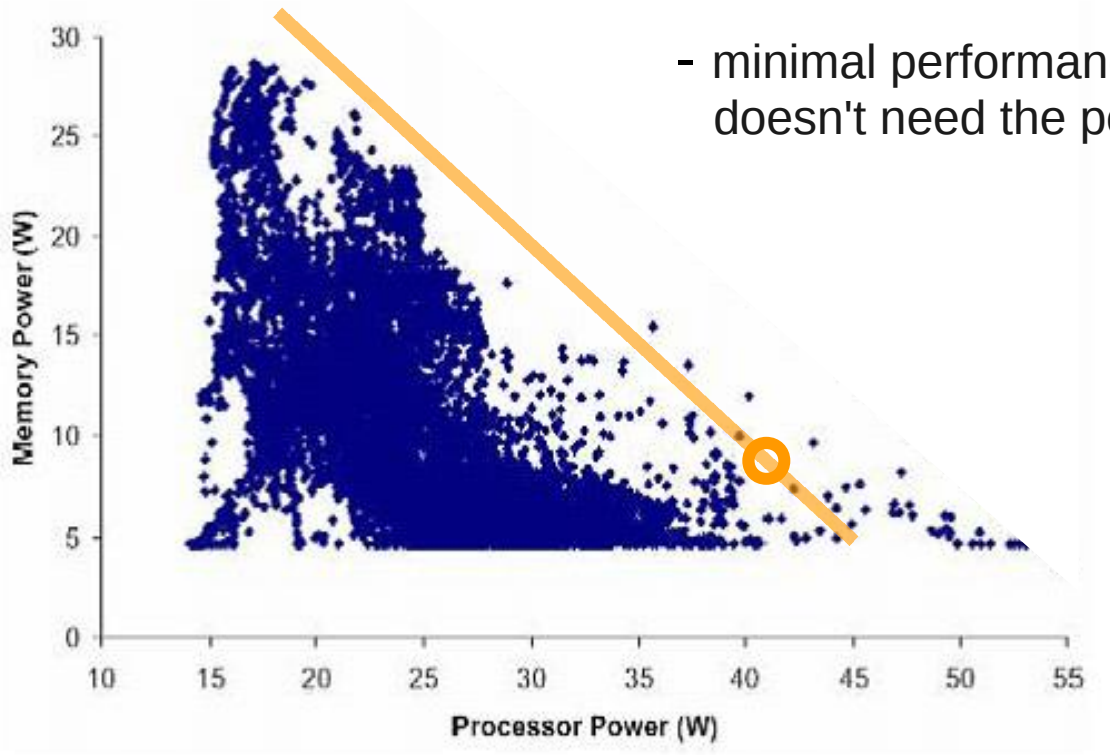

For tight total power budgets:
trim each sub-budget in proportion to previous consumption to preserve processor-memory ratio

Figure and observations from W. Felter, K. Rajamani, C. Rusu, and T. Keller,
"A Performance-Conserving Approach for Reducing Peak Power Consumption in Server Systems,"
in Proceedings of the 19th ACM Internaional Conference on Supercomputing, June 2005.

Each dot represents estimated power in a simulation snapshot during workload execution.

.....After

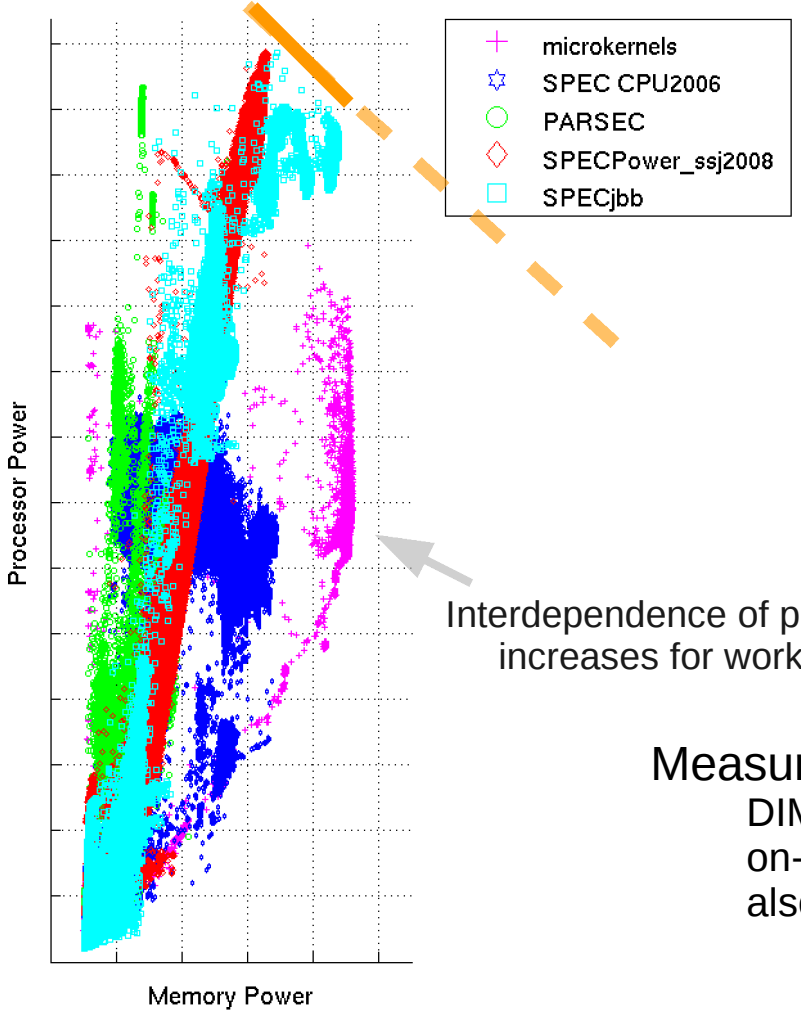

Multi-core, multi-threaded processors

4 mem controllers: more memory requests executing in parallel
128 threads: more core activity in parallel with memory accesses

What happens to power shifting opportunity?

- Processor & memory both active simultaneously
- Need to make intelligent budget choices
- Which component can make the best use of each Watt?

Performance doesn't scale 1:1 with power budget
Performance sensitivity to power budget varies through time and workload-to-workload

Interdependence of processor & memory power increases for workloads with large data footprints

Measured power:

DIMMs in memory power sensor
on-die memory controllers, etc. in processor power
also includes memory buffers (~constant) in proc. power

Each dot represents measured data from 32ms snapshot during workload execution on 32-core POWER7.
1 Watt has equal width and height on graph; absolute Watts not disclosed.

Experiments: Policies

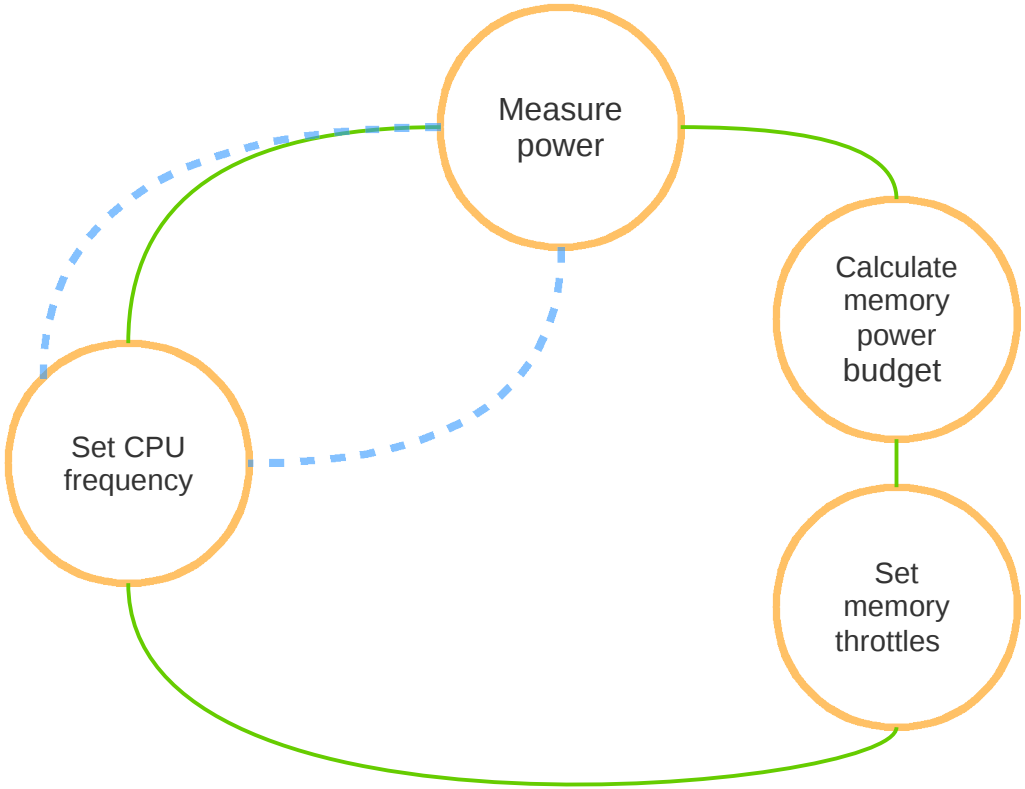

Simplified state machine diagram
(actual code has handshaking & safety features)

- **Total Power Cap**
  - Let memory run freely, manage total power budget with built-in DVFS power-cap loop
  - “Natural” shifting?
- **Proportional Budget**
  - Assign & enforce separate budgets
  - Policy from prior published work
- **How well do the policies handle multi-core systems?**

Experiments: Infrastructure

IBM POWER7 750 system

- 4 processors x 8 cores each x SMT 4 = 128 threads
- 64 GB, 1066 MHz DDR3 DIMMs
- Built-in power sensors & microcontroller for power/thermal management

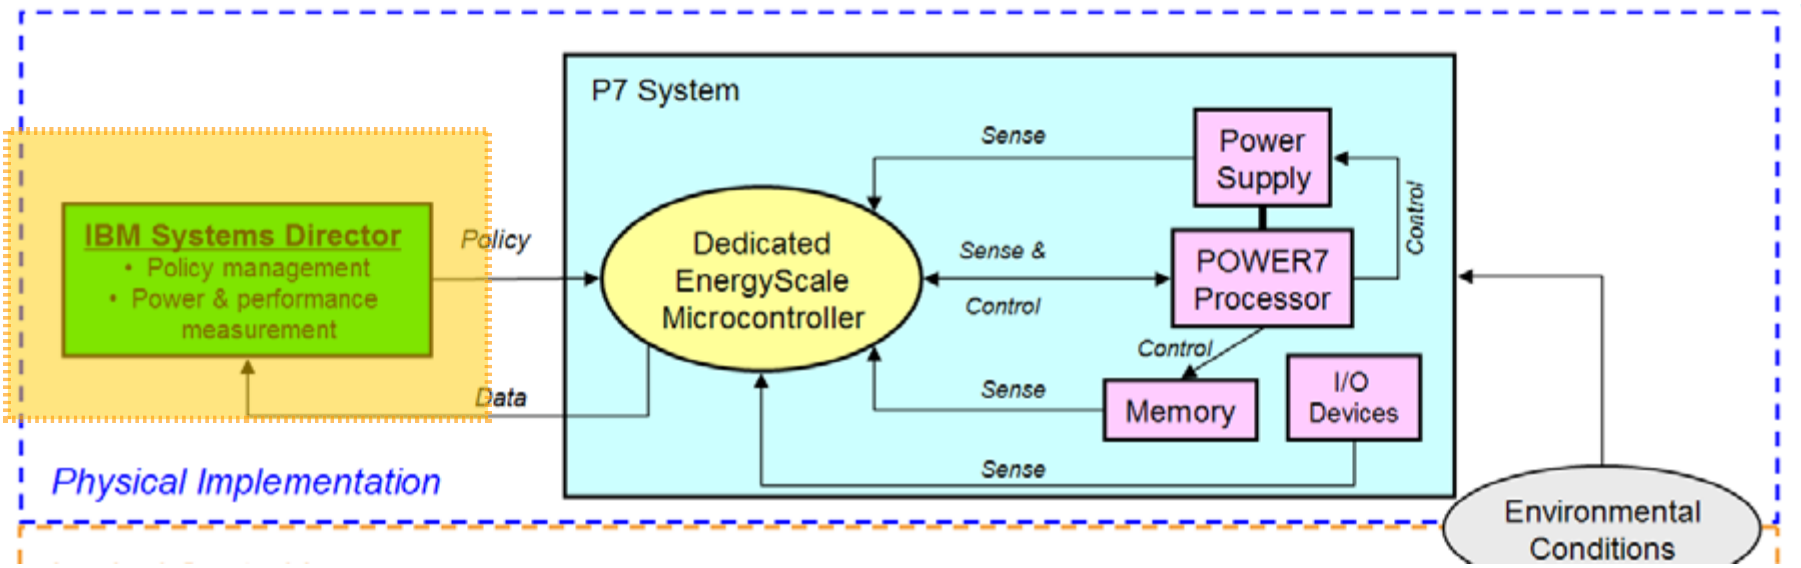

Figure from Floyd, et al., "Introducing the Adaptive Energy Management Features of the Power7 Chip," IEEE Micro, March/April 2011 (volume 31, number 2).

- System datasheet: http://www-03.ibm.com/systems/power/hardware/750

Experiments: Benchmarks

- SPEC CPU2006 int & fp [C, Fortran]
- SPECjbb2005 [Java; 3-tier client/server warehouse]
- SPECPower_ssj2008 [Java (based on jbb); intensity steps 100% down to idle]
- PARSEC [C; multi-threaded, shared-memory kernels]
- Custom microbenchmarks [C; target specific proc/mem points]
  - DAXPY (double A * X plus Y)
  - Random memory accesses
  - Floating-point multiply-accumulate
  - Sequences of square root operations
  - Dataset size for each is selectable, from L1-contained to main memory

Experiments: Results

- **Baseline:** high power limit (no management needed)
- **Moderate budget**
  - Total power limit near “typical” workload power draw
  - Outcome similar for both policies
    - Similar frequency choices, similar performance impact
- **Constrained budget**
  - Low total power limit requires active management
  - Surprising result: Proportional Budget policy chose slightly lower frequencies

    Investigation: allocated—but wasted—memory power budget reduced power available to processors

Discussion