

*Computation and Coordination
on the Information
Superhighway*

OR

**A Tale of Two Systems
Projects in URCS**

David Albonesi

Sandhya Dwarkadas

University of Rochester



A Tale of Two Projects

- **Complexity-Adaptive Processing (CAP)**
 - Ultra-low-power, yet high performance, self-tuning microprocessors that adapt to the changing needs of software
- **InterWeave**
 - Efficient task coordination across networks



Motivation for CAP: The Power Threat

- **Microprocessor power consumption is a serious and growing problem at all computer product levels**
 - **Portables: battery life is diminishing (3-6 hours typical in laptops)**
 - **PCs: heat driving up per-unit microprocessor costs (\$1 per each W over 40W)**
 - **Servers: power can be a significant cost (25%) of data centers**



Motivation for CAP: The Power Threat

- **Microprocessor power consumption is a serious and growing problem at all computer product levels**
 - **Portables: battery life is diminishing (3-6 hours typical in laptops)**
 - **PCs: heat driving up per-unit microprocessor costs (\$1 per each W over 40W)**
 - **Servers: power can be a significant cost (25%) of data centers**
- **At least one large corporation has instituted a “power scrutiny” policy for all future computer purchases**



Motivation for CAP: The Power Threat

- **Microprocessor power consumption is a serious and growing problem at all computer product levels**
 - **Portables: battery life is diminishing (3-6 hours typical in laptops)**
 - **PCs: heat driving up per-unit microprocessor costs (\$1 per each W over 40W)**
 - **Servers: power can be a significant cost (25%) of data centers**
- **At least one large corporation has instituted a “power scrutiny” policy for all future computer purchases**
- **Intel issued a “plea for help” at a recent industry conference**
 - **Fred Pollack (Director, Microprocessor Research Labs): Power is *the* limiter to being able to increase performance in the future**
 - **Deo Singh (Manager of server architecture): Power for servers may constitute 25% of the data center costs for an ISP**
 - **George Cai (Lead microarchitect): High frequency architectures must be power efficient as well as high performance**



Motivation for CAP: The Power Threat

- **Microprocessor power consumption is a serious and growing problem at all computer product levels**
 - **Portables: battery life is diminishing (3-6 hours typical in laptops)**
 - **PCs: heat driving up per-unit microprocessor costs (\$1 per each W over 40W)**
 - **Servers: power can be a significant cost (25%) of data centers**
- **At least one large corporation has instituted a “power scrutiny” policy for all future computer purchases**
- **Intel issued a “plea for help” at a recent industry conference**
 - **Fred Pollack (Director, Microprocessor Research Labs): Power is *the* limiter to being able to increase performance in the future**
 - **Deo Singh (Manager of server architecture): Power for servers may constitute 25% of the data center costs for an ISP**
 - **George Cai (Lead microarchitect): High frequency architectures must be power efficient as well as high performance**
- **Interest in Transmeta**



Motivation for CAP: The Power Threat

- **Microprocessor power consumption is a serious and growing problem at all computer product levels**
 - **Portables: battery life is diminishing (3-6 hours typical in laptops)**
 - **PCs: heat driving up per-unit microprocessor costs (\$1 per each W over 40W)**
 - **Servers: power can be a significant cost (25%) of data centers**
- **At least one large corporation has instituted a “power scrutiny” policy for all future computer purchases**
- **Intel issued a “plea for help” at a recent industry conference**
 - **Fred Pollack (Director, Microprocessor Research Labs): Power is *the* limiter to being able to increase performance in the future**
 - **Deo Singh (Manager of server architecture): Power for servers may constitute 25% of the data center costs for an ISP**
 - **George Cai (Lead microarchitect): High frequency architectures must be power efficient as well as high performance**
- **Interest in Transmeta**

*Bottom line: Power dissipation is perhaps
the most serious threat to future computing advances*



Microprocessor Power Trends

- **Desktop (Intel Pentium)**
 - 15W in 1992 to over 60W today
- **Server (Compaq Alpha)**
 - 30W in 1992 to over 90W today
- **By contrast, disk and memory power have remained fairly steady:**
 - IBM 36GB disk: 15W (about same as 2GB in 1994)
 - Micron 128MB memory DIMM: 5W (about same as 16MB in 1995)



What Impacts Microprocessor Power?

- Microprocessor dynamic power consumption

$$P \approx C_{\text{eff}} V_{\text{dd}}^2 f_{\text{clk}} N \alpha$$

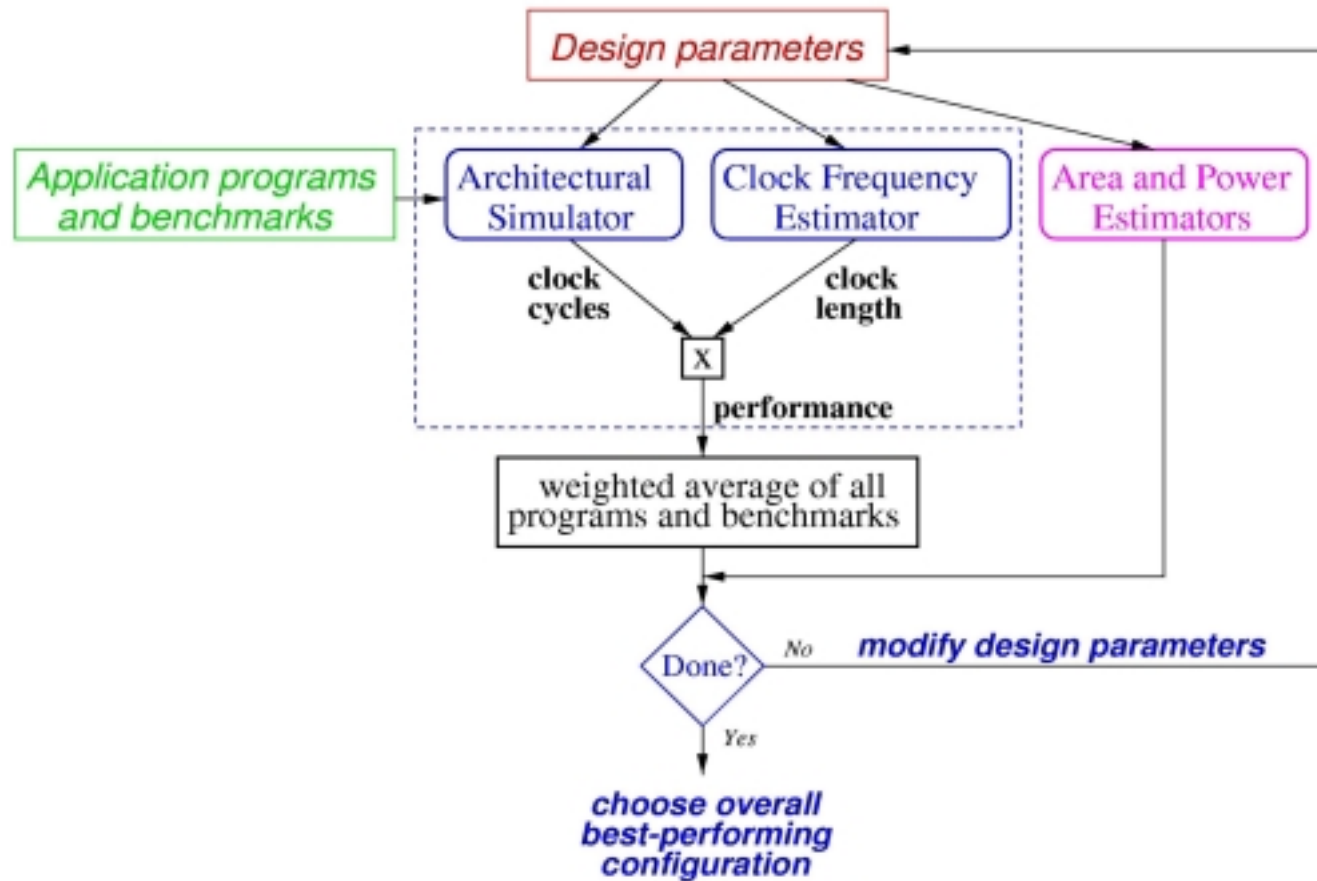
C_{eff} - loading V_{dd} - voltage f_{clk} - clock N - number of transistors α - %N switched

- A key point: the transistors, voltage, and clock frequency should be *judiciously applied* to the task at hand to maximize performance in the most power efficient manner possible




Microprocessor Design Today

- The usual approach (simplified view)



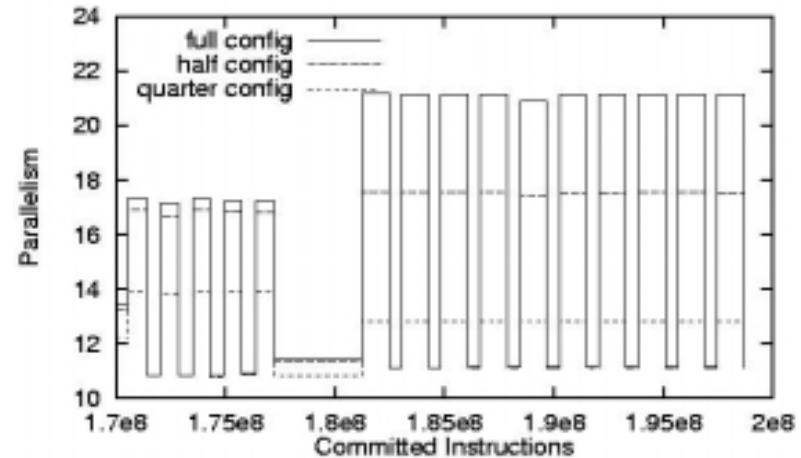
A Major Drawback: Inefficient Usage of Microprocessor Hardware Resources

- Hardware effectiveness may vary during the execution of a single program
 - Parallelism variation with issue and execution hardware complexity




A Major Drawback: Inefficient Usage of Microprocessor Hardware Resources

- Hardware effectiveness may vary during the execution of a single program
 - Parallelism variation with issue and execution hardware complexity



- Parallelism variation with data cache size

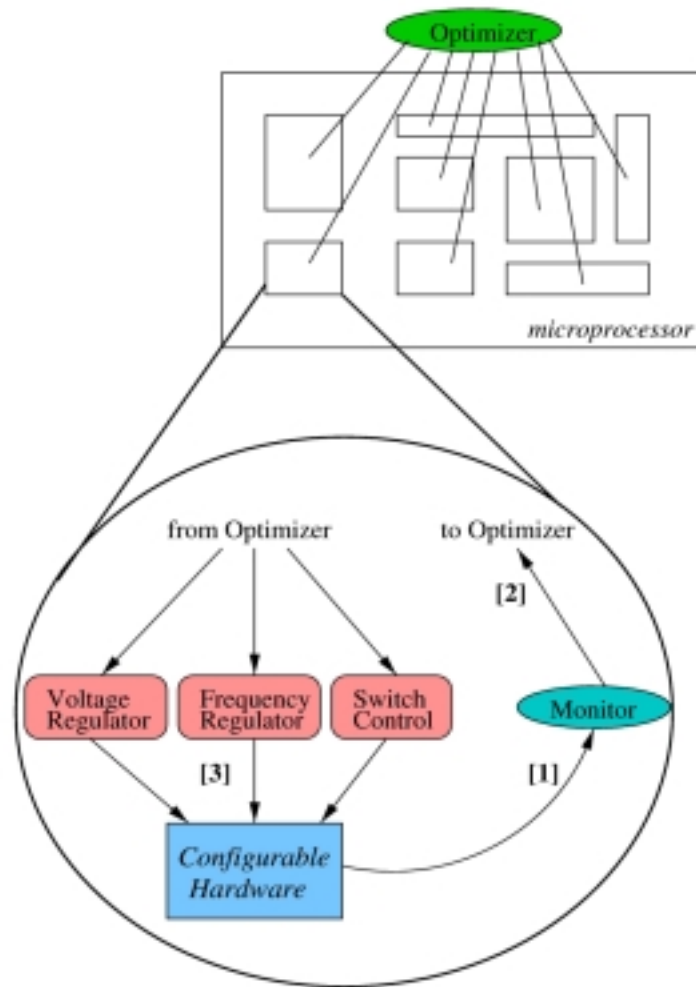


The CAP Approach

- ***Dynamically tune*** each component of power and performance to match changing software needs
 - The size, voltage, and clock frequency of each microprocessor hardware structure is *flexible*
 - Special circuitry *monitors* applications as they run
 - Special software *adjusts* the hardware structure size, frequency, and voltage to provide the required performance at the lowest power
- **The potential impact: a 5-10X reduction in power with no loss of performance**
- **Project covers the spectrum from circuits to operating systems**



CAP: A Self-Tuning Microprocessor



For each hardware unit:

[1] A *monitor* circuit detects a new program phase and determines how *efficiently* the *configurable hardware* is being used by the program and how *speed-critical* the hardware unit is

[2] The *optimizer* software determines what combination of voltage, frequency, and hardware configuration will provide enough performance at the lowest possible power for the next period of operation

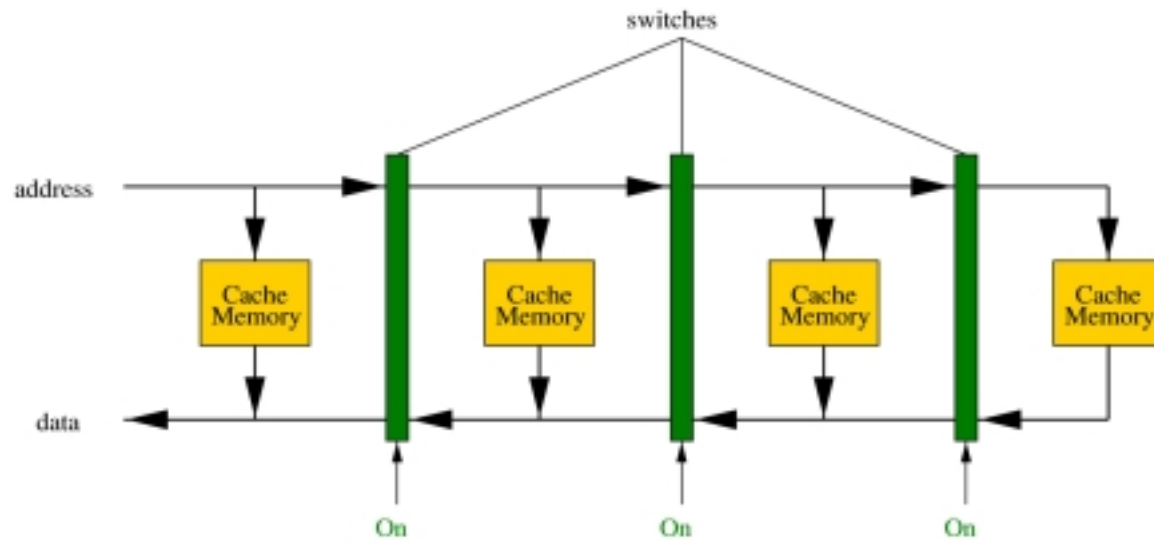
[3] These parameters are adjusted and the process continually repeats

Dynamic tuning of power and performance factors within each unit of the chip



Configurable Hardware

- Example of a configurable cache memory

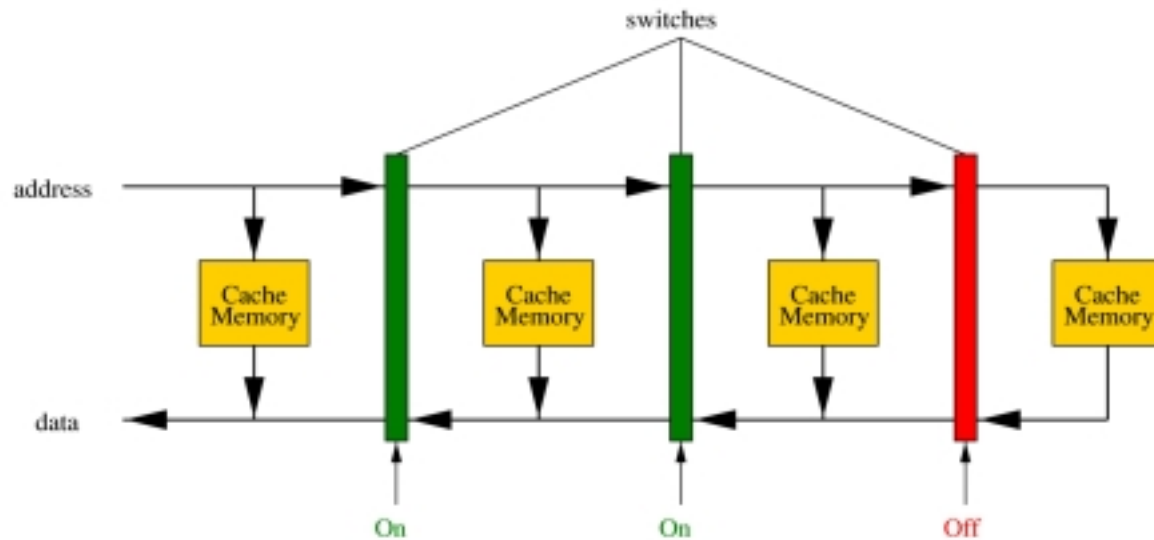


- Identical elements separated by switch circuits



Configurable Hardware

■ Example of a configurable cache memory




- Identical elements separated by switch circuits
- Switches can be opened to turn off one or more elements when they are not needed
 - Reduces amount of hardware charged each cycle => lower power
 - Reduces time to access data
 - Potential higher performance via increasing frequency OR
 - Even lower power via lowering the voltage to restore the original delay



Making *Fast* and *Dense* Configurable Hardware


- The key: exploit *memory partitioning* and *wire repeaters*

Then: single long address wire




Making *Fast* and *Dense* Configurable Hardware

- The key: exploit memory *partitioning* and *wire repeaters*

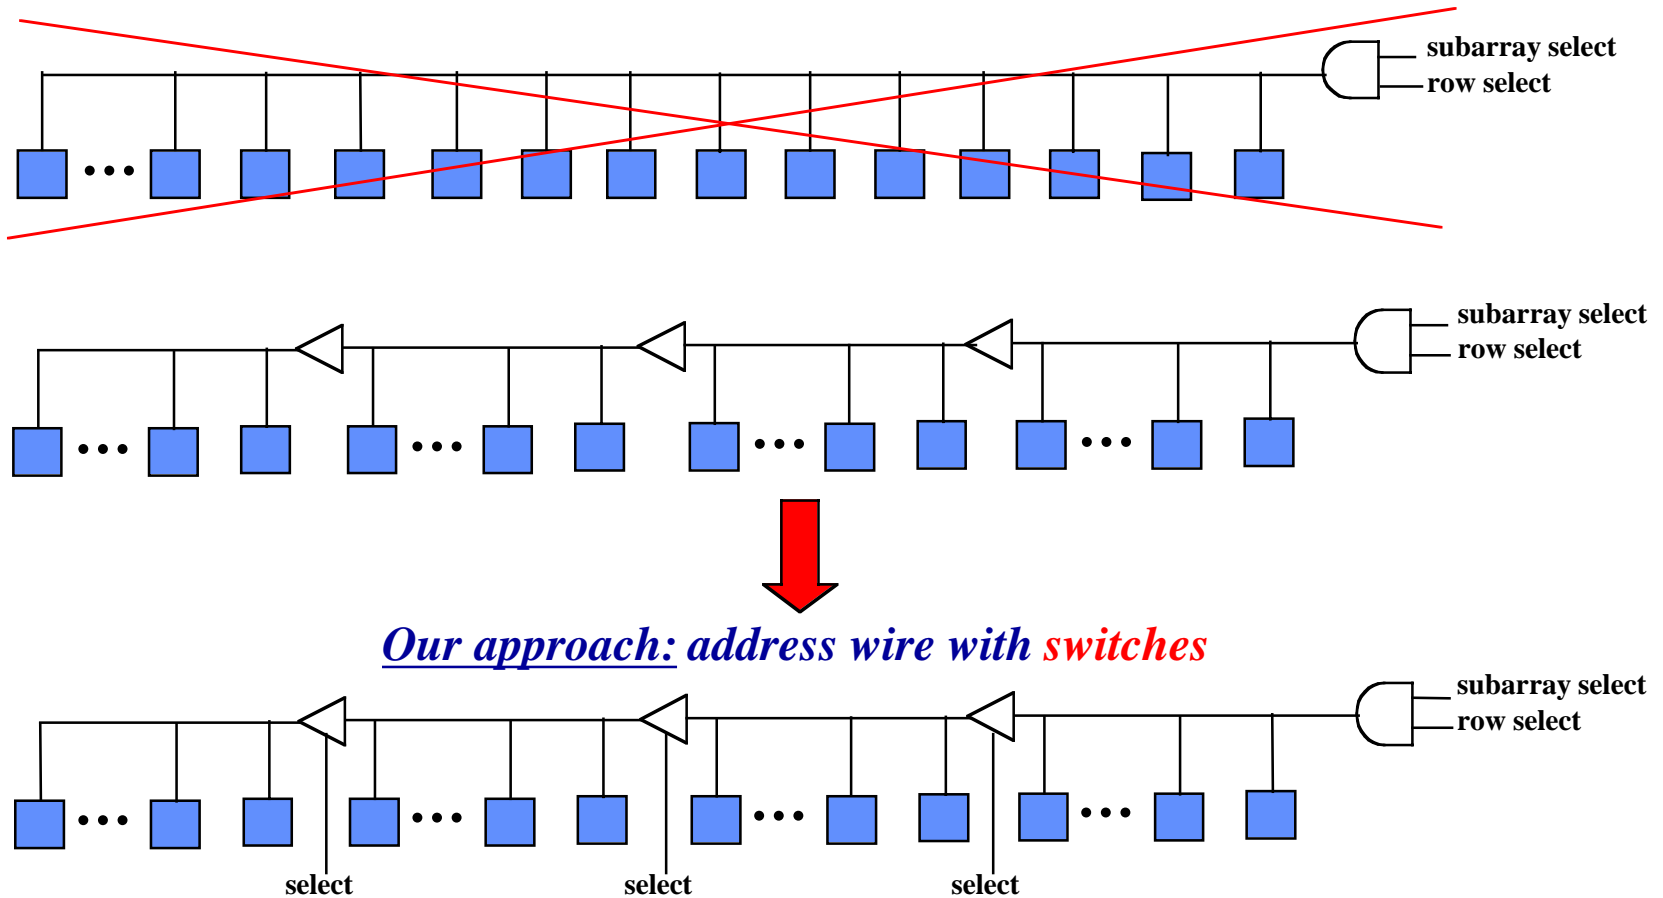


Now: address wire with repeaters for speed



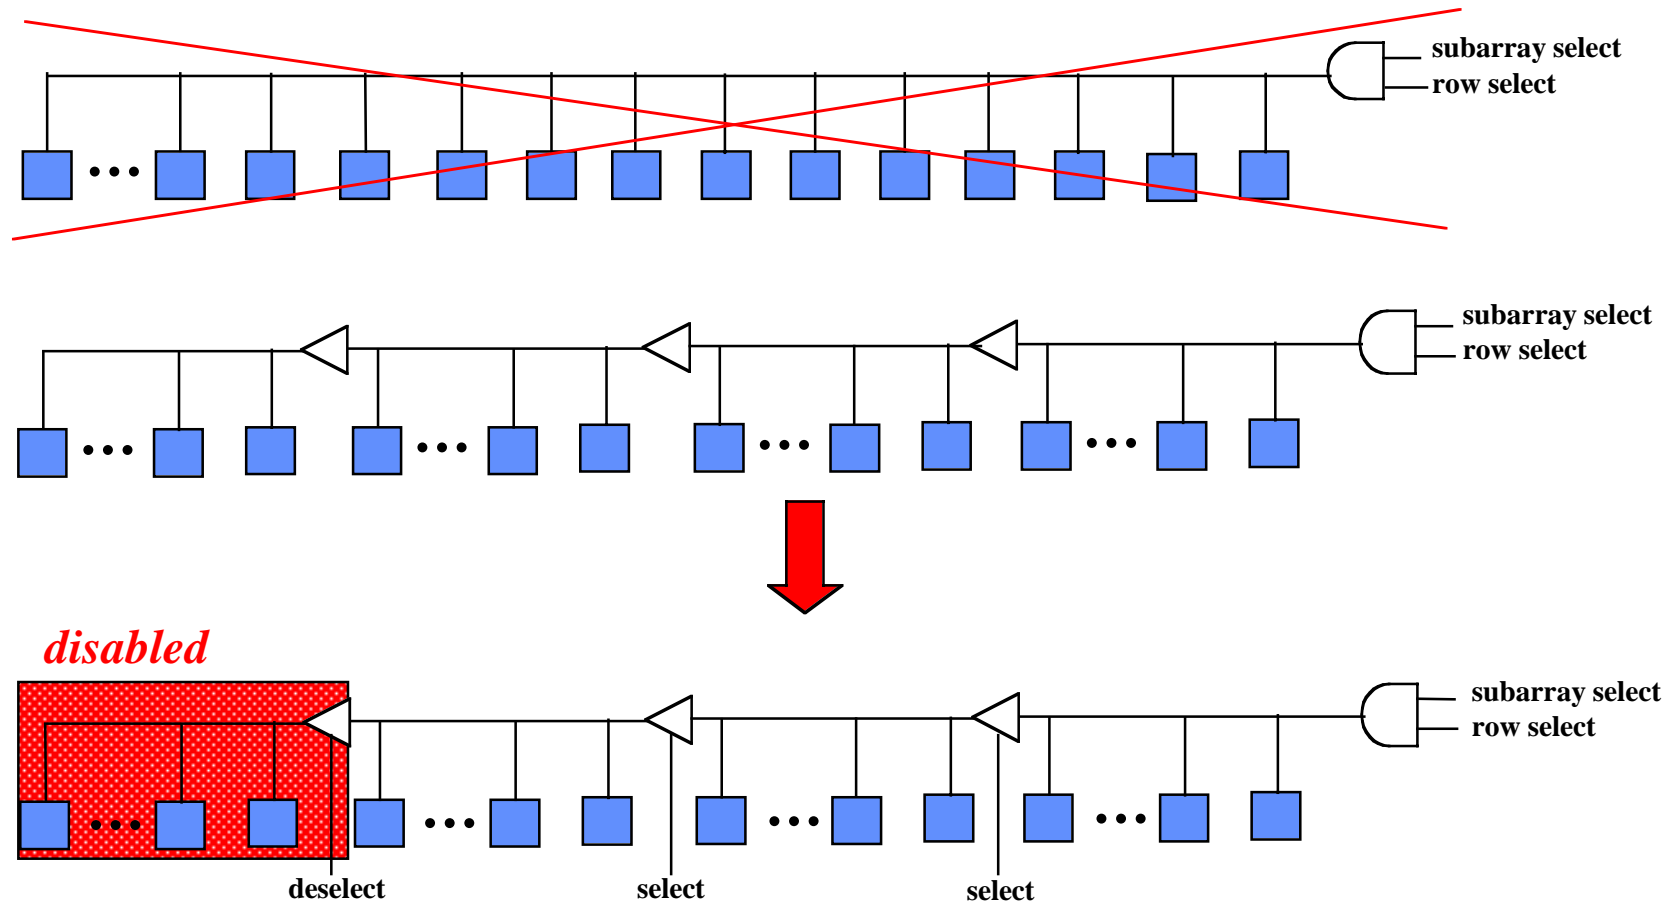
Making *Fast* and *Dense* Configurable Hardware

- The key: exploit *memory partitioning* and *wire repeaters*



Making *Fast* and *Dense* Configurable Hardware

- The key: exploit *memory partitioning* and *wire repeaters*

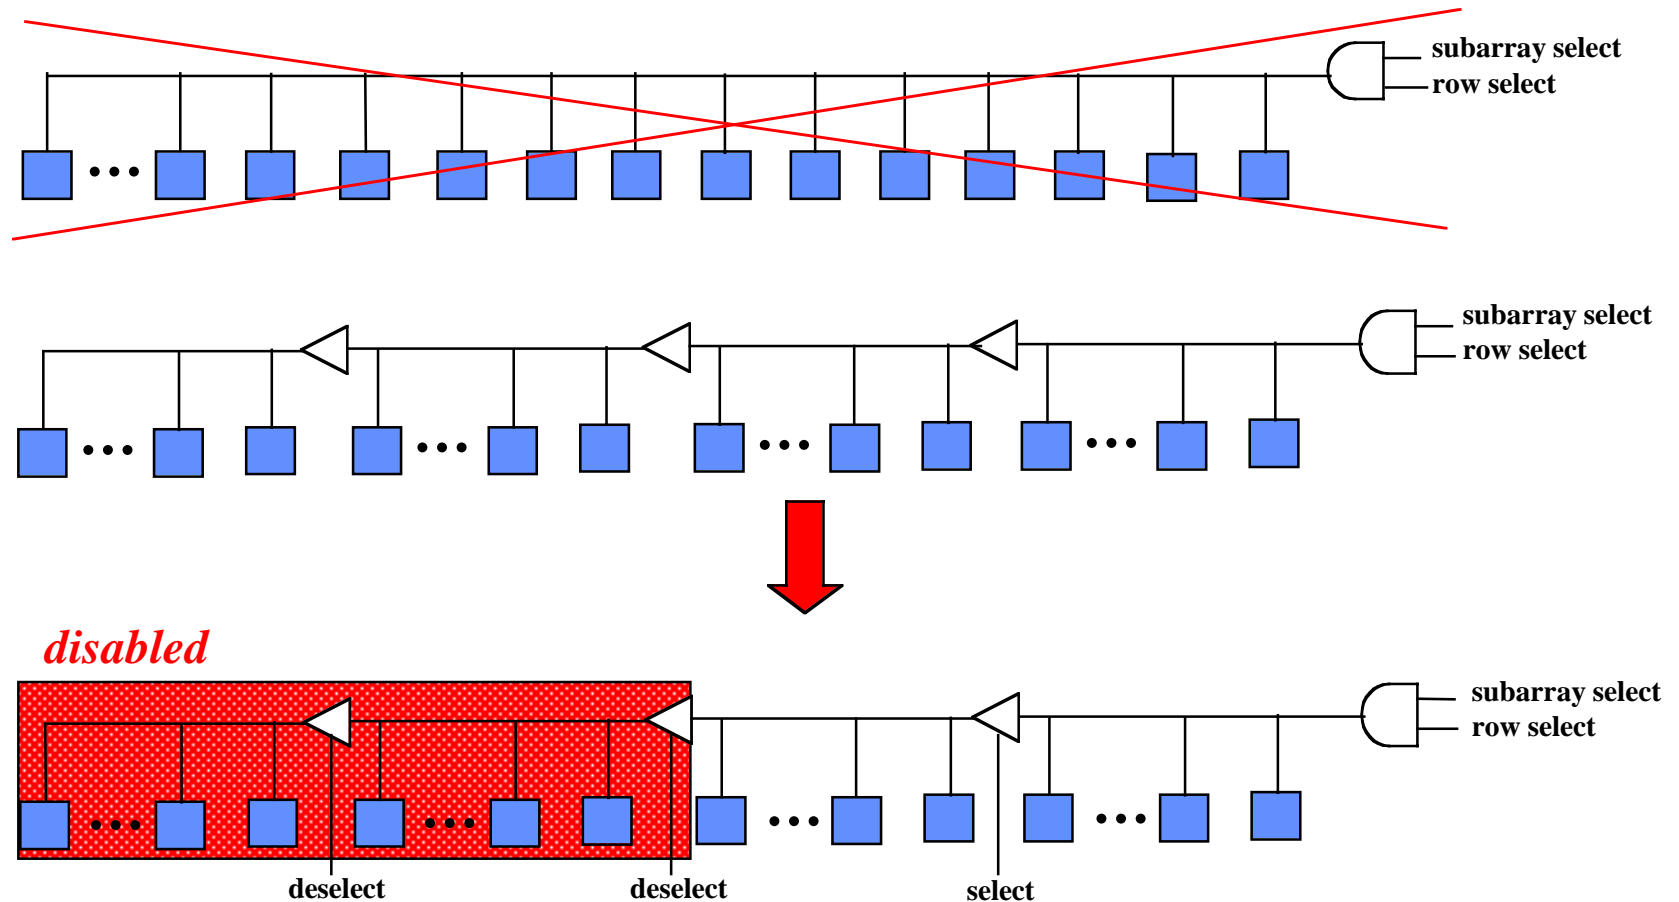


Hardware configurability with negligible speed and area impact



Making *Fast* and *Dense* Configurable Hardware

- The key: exploit *memory partitioning* and *wire repeaters*

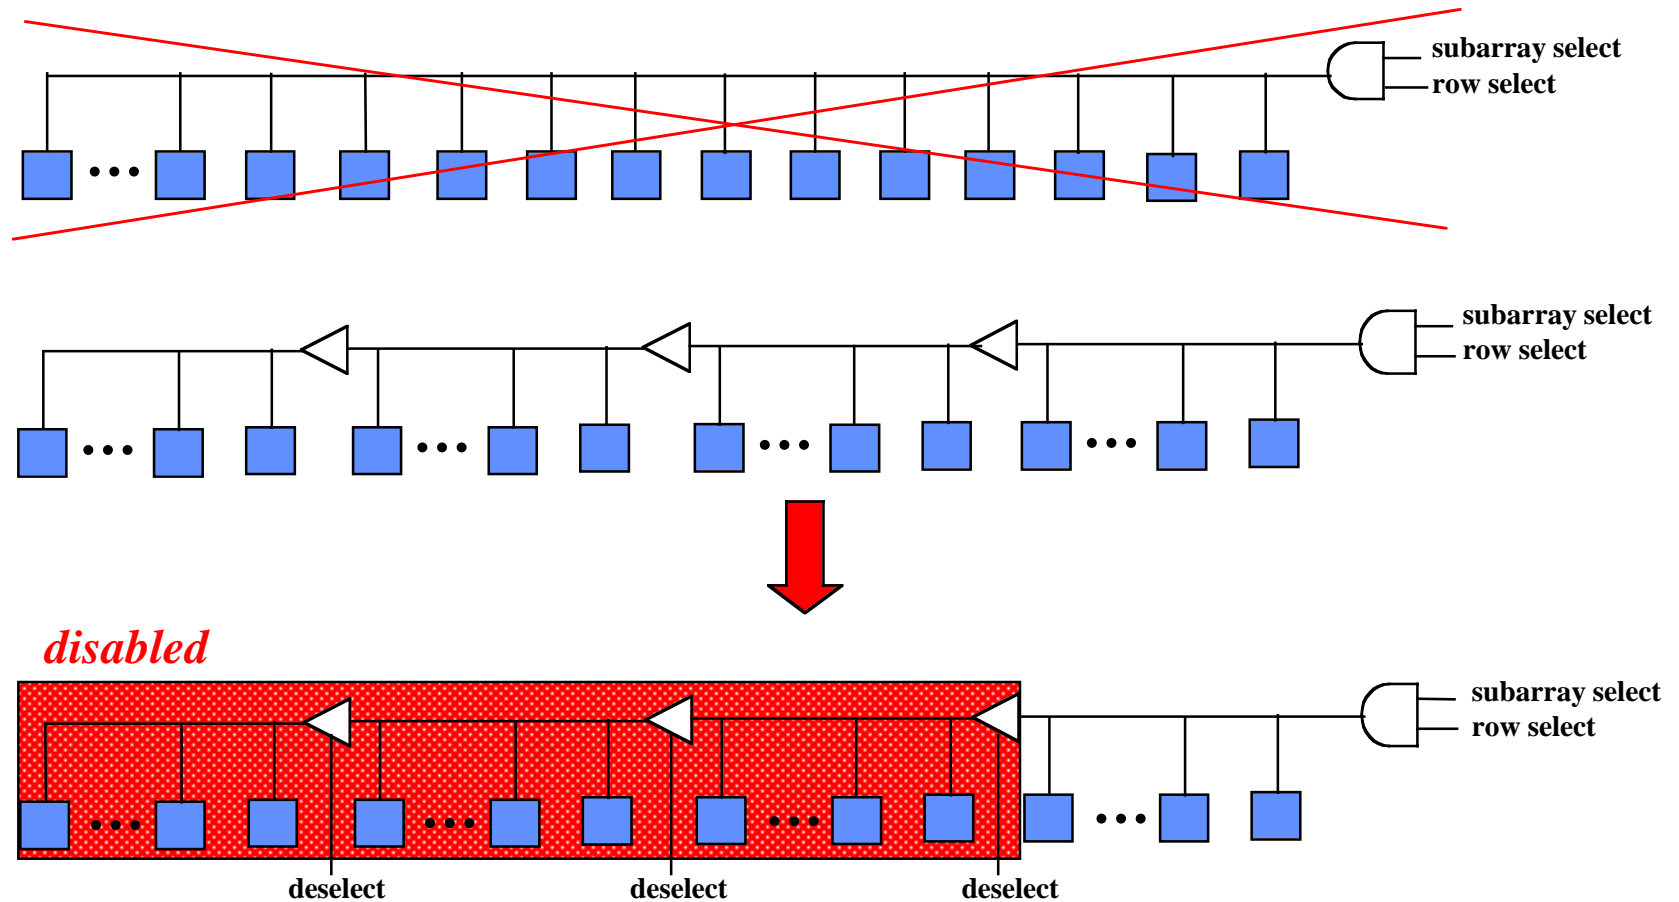


Hardware configurability with negligible speed and area impact



Making *Fast* and *Dense* Configurable Hardware

- The key: exploit *memory partitioning* and *wire repeaters*

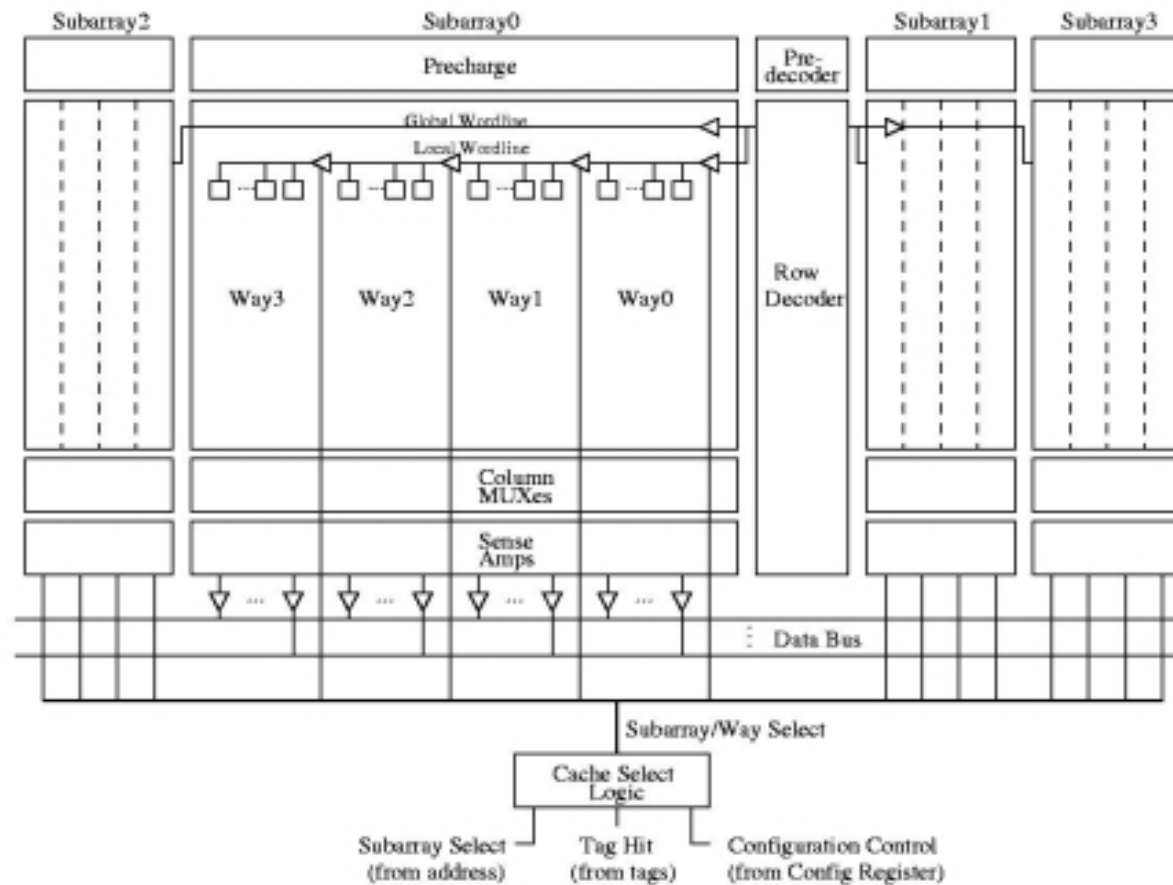


Hardware configurability with negligible speed and area impact



CAP Data Cache Hierarchy

- Leverages conventional 4-way set associative 2MB cache
- Cache ways and subarrays dynamically enabled/disabled at runtime



CAP Data Cache Hierarchy Configurations

- Eight possible L1/L2 configurations in one cache design
 - Vary in energy, size, associativity, and access latency



Subarray/Way Allocation (L1 or L2)

Cache Configuration	L1 Size	L1 Assoc	L1 Acc Time	Subarray/Way Allocation (L1 or L2)															
				Subarray 2				Subarray 0				Subarray 1				Subarray 3			
				W3	W2	W1	W0	W3	W2	W1	W0	W0	W1	W2	W3	W0	W1	W2	W3
256-1	256KB	1 way	2.0	L2	L2	L2	L2	L2	L2	L2	L1	L1	L2	L2	L2	L2	L2	L2	L2
512-2	512KB	2 way	2.5	L2	L2	L2	L2	L2	L2	L1	L1	L1	L1	L2	L2	L2	L2	L2	L2
768-3	768KB	3 way	2.5	L2	L2	L2	L2	L2	L1	L1	L1	L1	L1	L1	L2	L2	L2	L2	L2
1024-4	1024KB	4 way	3.0	L2	L2	L2	L2	L1	L1	L1	L1	L1	L1	L1	L1	L2	L2	L2	L2
512-1	512KB	1 way	3.0	L2	L2	L2	L1	L2	L2	L2	L1	L1	L2	L2	L2	L1	L2	L2	L2
1024-2	1024KB	2 way	3.5	L2	L2	L1	L1	L2	L2	L1	L1	L1	L1	L2	L2	L1	L1	L2	L2
1536-3	1536KB	3 way	4.0	L2	L1	L1	L1	L2	L1	L1	L1	L1	L1	L1	L2	L1	L1	L1	L2
2048-4	2048KB	4 way	4.5	L1	L1	L1	L1	L1	L1	L1	L1	L1	L1	L1	L1	L1	L1	L1	L1



CAP Data Cache Hierarchy Results

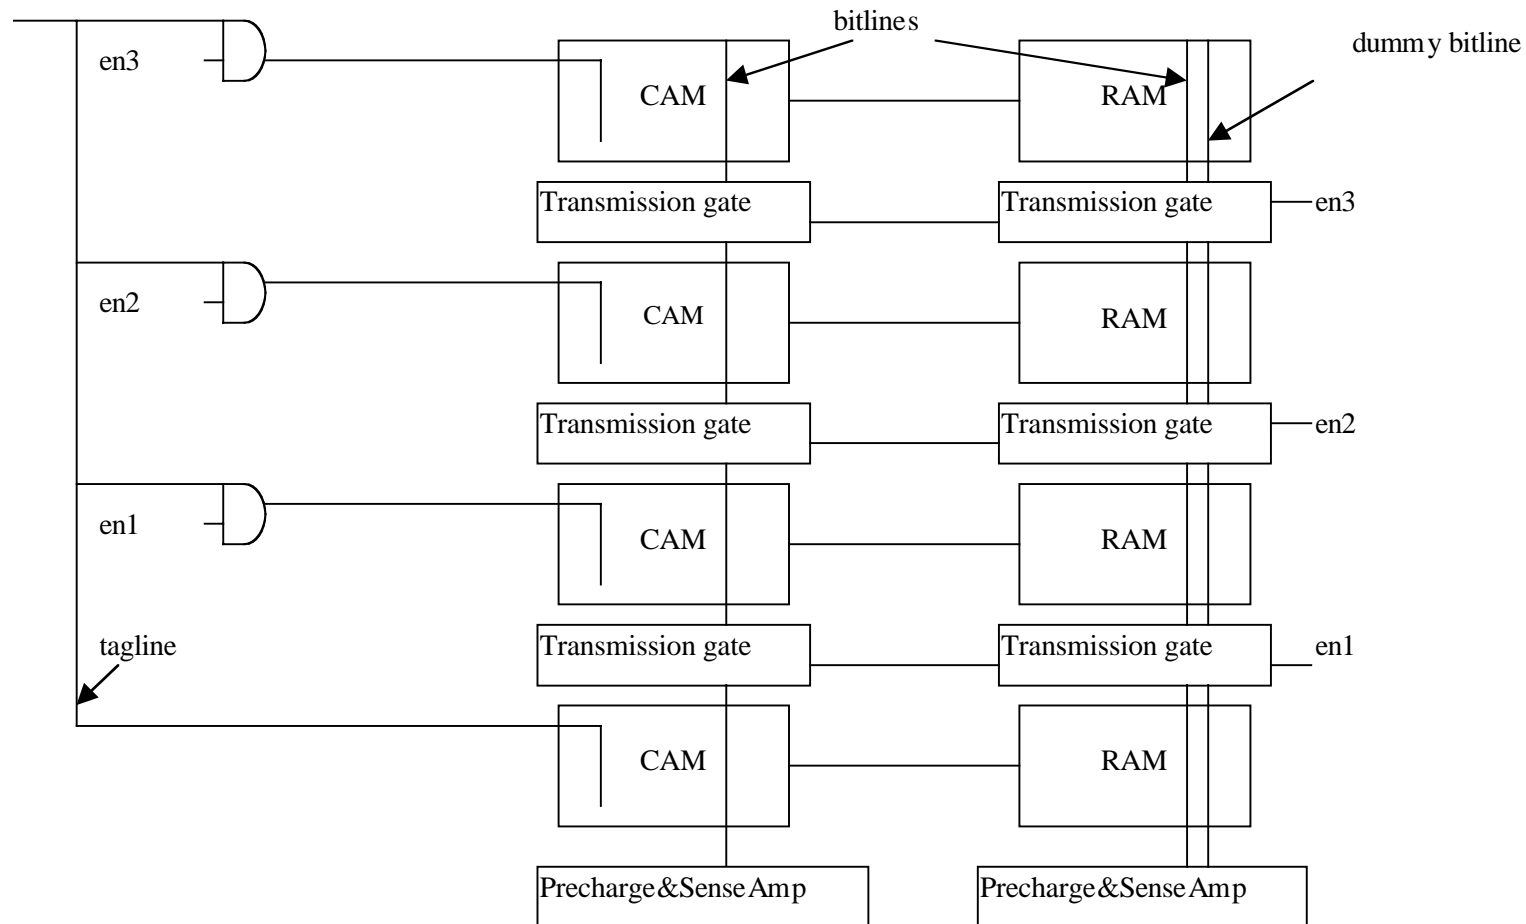
- Comparison of conventional 3-level and dynamic memory hierarchies
 - Fixed L1 cache and configurable L2 (serial tag-data access for all L2's)



- Adaptive design matches best performance while dissipating considerably less energy (up to 2.5X) than highly optimized conventional hierarchy
 - Only addressing the switched capacitance power component



CAP Issue Queue



- Shutdown disabled 8-entry “chunks” based on usage
- Reduces queue power by up to 3.5X
- No clock cycle time impact




IBM Collaboration

- **Designing two low-power test chips in collaboration with IBM**
 - Test combination of CAP and IBM techniques
 - Fabricated in advanced IBM process
 - Demonstrate very high clock speed yet extremely low power
- **Phase I: Fabricate and test CAP issue queue and functional unit**
 - Target tape-out date: end 2001
- **Phase II: PowerPC-based processor incorporating CAP and IBM techniques in several units**
 - Perhaps based on Power-3 or Power-4 design
 - Target tape-out date: mid to end 2002
- **Intel also interested**
 - Speaking at internal Intel Low-Power Research Symposium in Santa Clara on November 10th



The CAP Team

■ Faculty



□ l-r: Dave Albonesi, Michael Scott, Eby Friedman, Sandhya Dwarkadas

■ Graduate Students

- Rajeev Balasubramonian
- Alper Buyuktosunoglu
- Volkan Kursun
- Grigorios Magklis
- Greg Semeraro

■ Staff

- Keith Parkins
- Jim Roche
- Greg Semeraro

■ Two nostdocs to come

