

---

# Optimizing the Thermal Behavior of Subarrayed Data Caches

---


Johnsy K. John, **Jie Hu**, and Sotirios G. Zivarras  
Electrical and Computer Engineering  
New Jersey Institute of Technology

---

## Outline

- Motivation
- Introduction
  - HotSpot temperature model
  - Cacti cache subarray model
- Potential thermal problems in subarrayed caches
- Thermal-aware cache subarraying
  - Separated cache subarraying
  - Way-interleaved cache subarraying
  - Interaction between thermal control and leakage control
- Experiments and results
- Conclusions

## Power: A Design Limiter



Power density too high to keep junctions at low temperature

S. Borkar, Intel  
ICCD 2005, San Jose

10/05/2005

3

## Thermal Management

- Adversities due to high chip temperature
  - Dramatically increases costs of cooling systems and chip packaging
  - Causes exponential increase in leakage
  - Lowers chip reliability
- Dynamic thermal management (DTM)
  - Pre-defined DTM trigger temperature
  - Dynamically monitors chip-wide temperature
  - Invokes power reduction schemes (DVS/DFS, clock gating, speculation controlling, etc.) as response to avoid thermal emergencies
- Static thermal-aware microarchitectures
  - DTM incurs performance overhead once response is triggered
  - Design thermal-aware microarchitectural components to reduce the offense of exceeding DTM trigger temperature thus reducing the performance overhead
  - Combined to improve the efficiency of DTM

10/05/2005

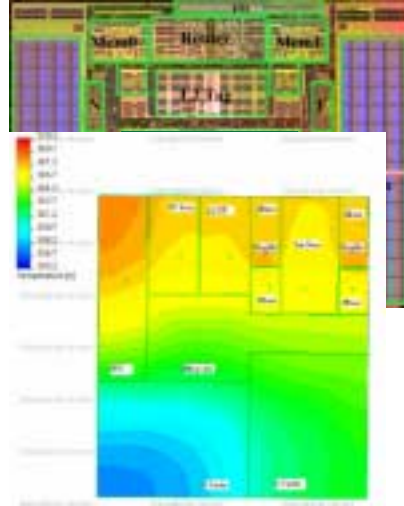
ICCD 2005, San Jose

4

## Why Caches, Not in Hot Spot Yet?

- Alpha 21364 Technology
- 0.18  $\mu\text{m}$  CMOS
- 1250 MHz
- 135 Watts @ 1.65 volts
- 4 cm<sup>2</sup>
- 7 Layer Metal
- 152 million transistors
  - 15 million logic
  - 137 million SRAM

Data and die photo from Peter Bannon's talk  
 "Alpha 21364 (EV7)", 2002  
 Thermal picture from HotSpot ISCA'03  
 presentation "Temperature-Aware  
 Microarchitecture"




10/05/2005

ICCD 2005, San Jose

5

## Sub-threshold Leakage vs. Temperature




10/05/2005

ICCD 2005, San Jose

6

## HotSpot Thermal Model

- Stepwise temperature update in HotSpot [Skadron et. al.]  
 $\Delta T = P/C*\Delta t + T/RC*\Delta t$
- Heat diffusing/spreading among adjacent blocks: lateral resistances



- Important observations
  - Localized heating occurs fast
  - Floorplanning: hot components should be separated from each other to avoid heat buildup

10/05/2005

ICCD 2005, San Jose

7

## Thermal Behavior Inside Caches

- Understanding the thermal behavior of caches is the key to design thermal-aware caches
- Data localities in cache accesses indicate very different cache-wide activity/power/temperature
- Modeling caches as monolithic blocks (conventional way) cannot capture this feature
- Cache level thermal behavior is not sufficient and may not provide insight for cache thermal optimization
- **Question: how to characterize the cache thermal behavior and at what granularity?**

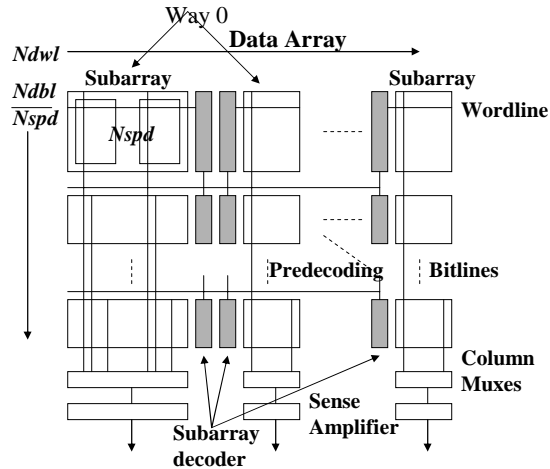
10/05/2005

ICCD 2005, San Jose

8

## Cacti Model for Subarrayed Caches

- In Cacti model, data array is subarrayed by  $N_{dwl}$ ,  $N_{dbl}$ , and  $N_{spd}$  to minimize an objective function of access latency and energy consumption
- $N_{dwl}$  defines how wordline is segmented horizontally
- $N_{dbl}$  defines how bitline is segmented vertically
- $N_{spd}$  defines how many sets mapped to a single wordline
- → **Subarray based thermal profiling**



10/05/2005

ICCD 2005, San Jose

9

## Experimental Setup

- Extended SimpleScalar to model Alpha 21364 microprocessor

Int/FP issue queue	20/15 entries
Ld/St queue	64
Active list	80
Int/FP Reg. file	80/72
Fetch/Decode/Commit width	4 inst. Per cycle
Int/FP issue width	4/2 inst. Per cycle
Function units	4 IALU, 2IMULT/IDIV, 2FALU, 1FMULT/FDIV/FSQRT, 2 Mem Ports
Branch Predictor	Tournament predictor
L1 ICache/DCache	<b>64KB, 2 ways, 64B blocks, 2 cycles</b>
L2 UCache	4MB, 8 ways, 128B blocks, 12 cycles
Memory I/DTLB	225 cycles first chunk, 12 cycles rest 128 entries, fully assoc.

10/05/2005

ICCD 2005, San Jose

10

## Power and Temperature Models

- Dynamic power model derived from Wattch [Brooks and Martonosi]
- HotLeakage [Zhang et. al.] for sub-threshold leakage and gate leakage calculation
- HotSpot for transient and steady state temperatures

HotSpot Parameters		
Technology	130nm	70nm
Clock Frequency	3GHz	5.6GHz
Supply Voltage	1.5V	1.0V
Ambient Air Temp.	45°C	
Package Thermal Resistance	0.8K/W	
Die	0.5mm thick, 15.9mmX15.9mm	
Heat Spreader	Copper, 1mm thick, 3cmX3cm	
Heat Sink	Copper, 7mm thick, 6cmX6cm	

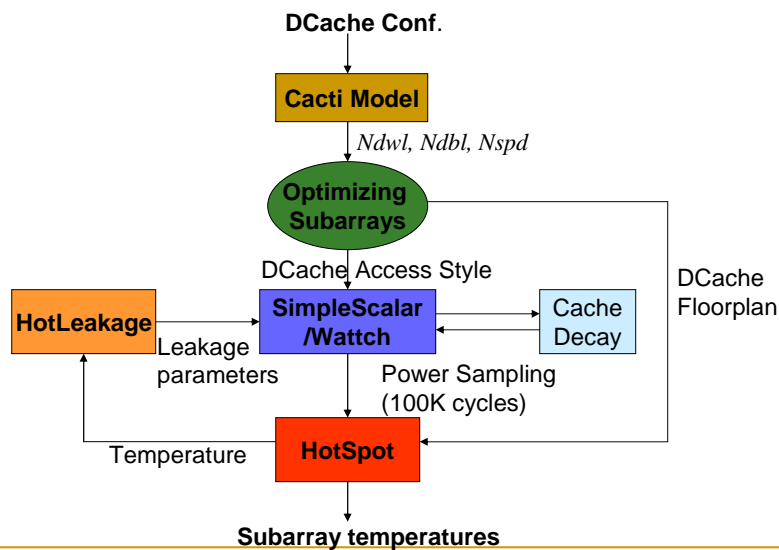
- Block sizes in floorplan scaled from 180nm to 130nm and 70nm

10/05/2005

ICCD 2005, San Jose

11

## Evaluation Framework



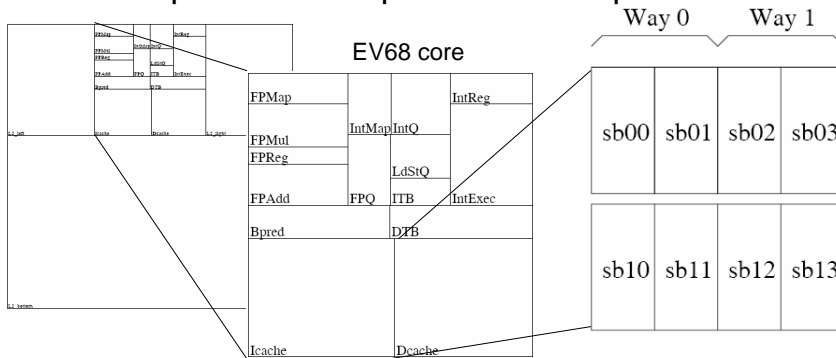
10/05/2005

ICCD 2005, San Jose

12

# Alpha 21364 Microprocessor Floorplan

- Default Alpha EV7 floorplan from HotSpot



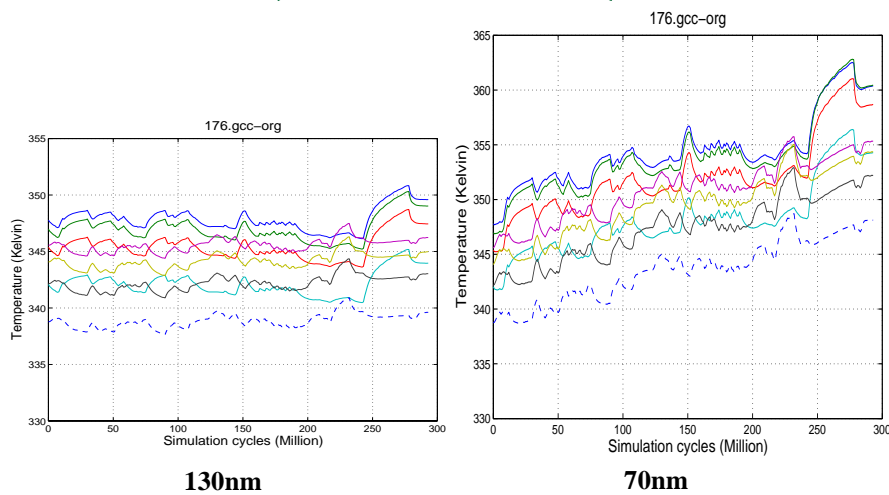
- Subarray layout from Cacti:  $Ndwl=4$ ,  $Ndbl=2$ ,  $Nspd=1$

10/05/2005

ICCD 2005, San Jose

13

# Thermal Behavior in Simple-Subarrayed Data Cache (130nm vs. 70nm)




10/05/2005

ICCD 2005, San Jose

14

## Simple Subarray Scheme

- Each read/write accesses all subarrays in a row
- Each accessed subarray can have at most 2 idle neighboring subarrays (above and below itself)
- Heat might build up in an accessed row due to the hot neighbors
- Data locality may continuously heat up a row
- High temperature leads significantly increased leakage power
- High leakage in turn contributes to each higher temperature, possibly leading to thermal runaway
- How to improve heat spreading/diffusing from hot subarrays to low-temperature subarrays?**



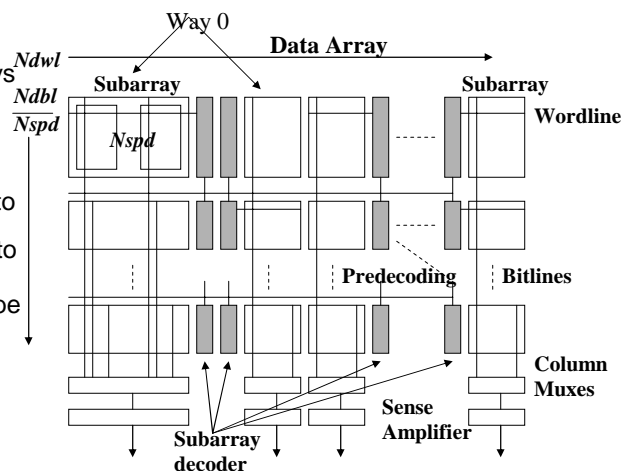
10/05/2005

ICCD 2005, San Jose

15

## Separated Cache Subarraying

- Possible solution: increase idle subarrays surrounding an accessed subarray
- Separated subarraying:** Subarrays belonging to a logical row are physically distributed to multiple rows
- Each subarray could be in one of  $N_{dbl}/N_{spd}$  rows within the same column
- Each subarray could have up to **4 idle subarrays**

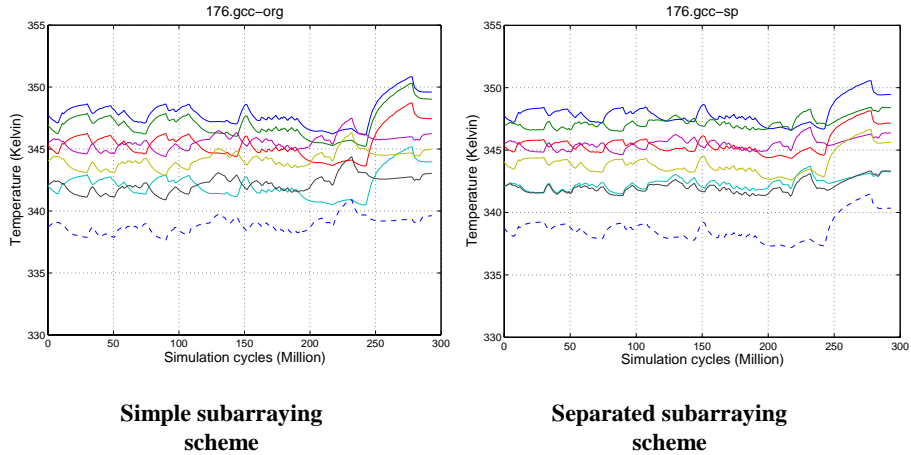


10/05/2005

ICCD 2005, San Jose

16

## Thermal Behavior in Separate-Subarrayed Data Cache (130nm)




10/05/2005

ICCD 2005, San Jose

17

## Separated Subarraying Scheme

- Limiting factors:
  - Data accesses exhibiting no tight locality may access multiple rows of subarrays
  - $Ndb1/Nspd$  limits the distance of separation of logically neighboring subarrays
  - In simulated Alpha 21364 L1 data cache, only 2 logic rows of subarrays
- Alternatives to increase idle subarrays surrounding an accessed subarray?



10/05/2005

ICCD 2005, San Jose

18

## Way-Interleaved Cache Subarraying

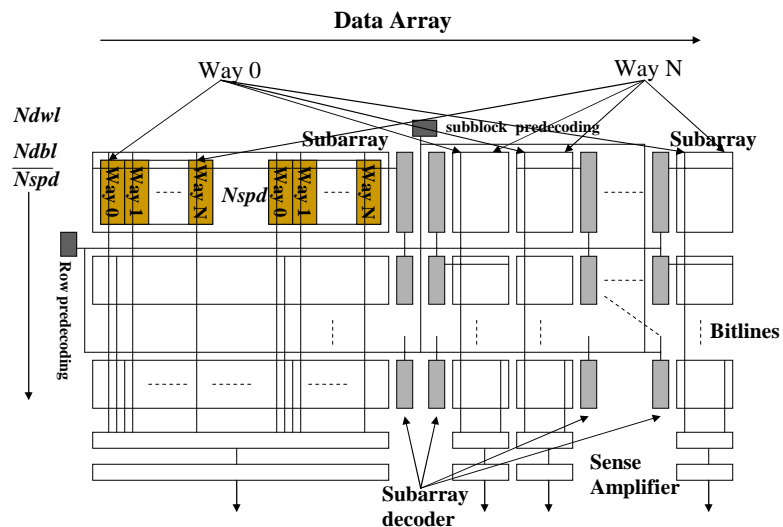
- Alternative approach to increase surrounding idle subarrays
  - Limit each access to only **ONE** subarray
- Distribute each way over subarrays in a row
- Each subarray holds a subblock from each of N ways
- Sub-block size =  $B/N_{dwl}$ 
  - Example, for the Alpha 21364 data cache,  $B=64$  bytes,  $N_{dwl}=4$ , subblock size =  $64/4 = 16$  bytes
- Row and subblock predecoding signals locate a particular subarray for access
  - Row predecoding uses higher  $\log_2(N_{dbl}/N_{spd})$  index bits
  - Subblock predecoding uses higher  $\log_2(N_{dwl})$  block offset bits
- **Significantly reduced dynamic power consumption!**

10/05/2005

ICCD 2005, San Jose

19

## Way-Interleaved Cache Subarraying

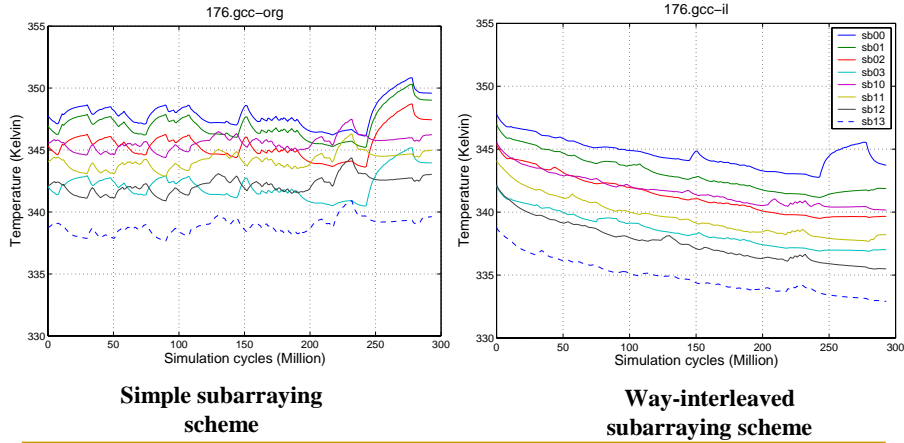


10/05/2005

ICCD 2005, San Jose

20

## Thermal Behavior in Interleaved-Subarrayed (130nm)

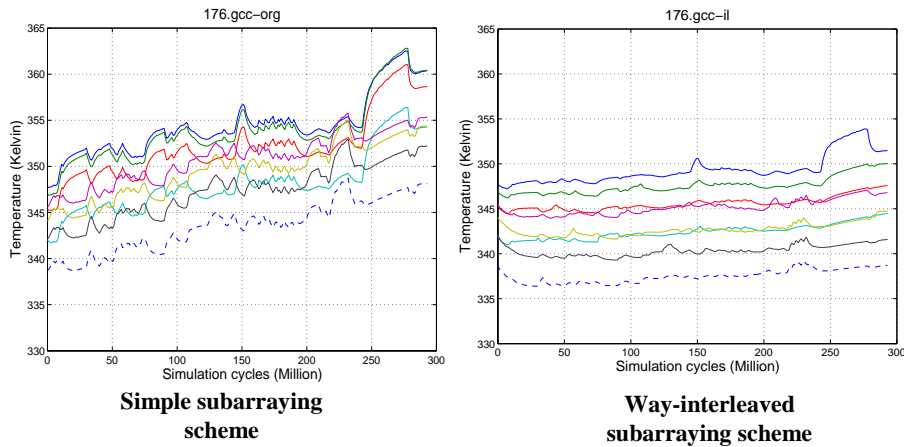


10/05/2005

ICCD 2005, San Jose

21

## Thermal Behavior in Interleaved-Subarrayed (70nm)

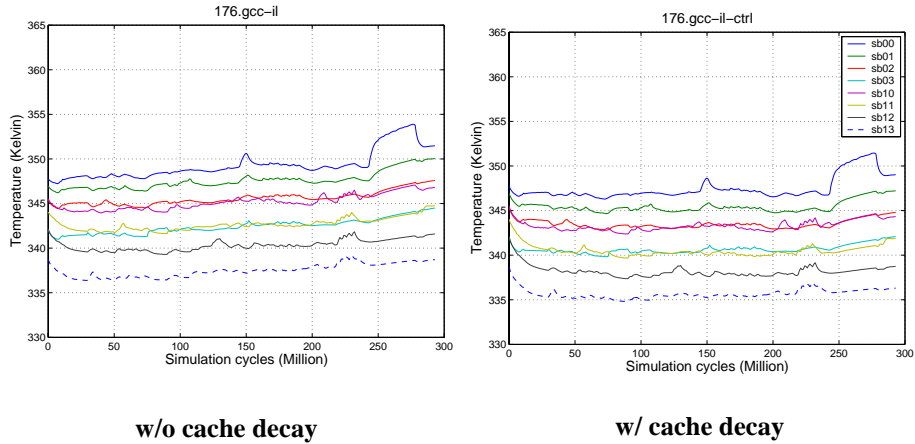


10/05/2005

ICCD 2005, San Jose

22

## Interleaved Subarraying with Cache Decay

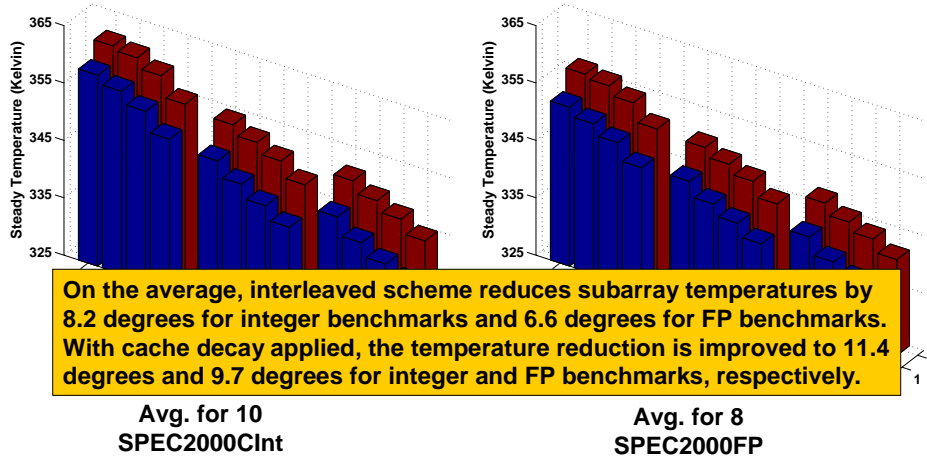


10/05/2005

ICCD 2005, San Jose

23

## Steady-State Temperature Reduction (70nm)

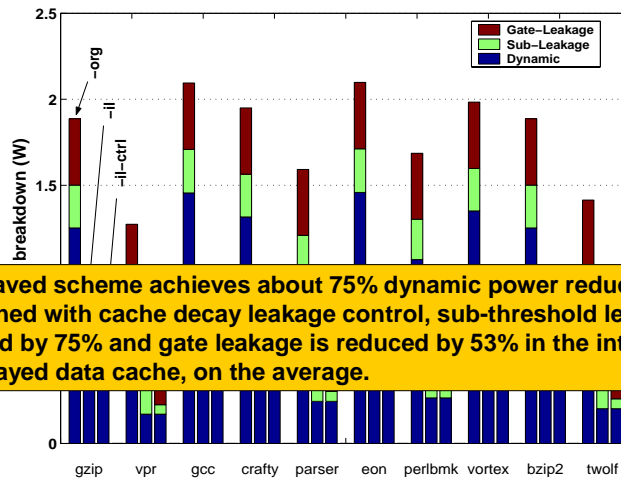


10/05/2005

ICCD 2005, San Jose

24

## Data Cache Power Breakdown (70nm)



**Interleaved scheme achieves about 75% dynamic power reduction. Combined with cache decay leakage control, sub-threshold leakage is reduced by 75% and gate leakage is reduced by 53% in the interleaved-subarrayed data cache, on the average.**

10/05/2005

ICCD 2005, San Jose

25

## Conclusions

- Investigated the thermal behavior inside subarrayed data cache
  - Subarrays exhibit quite different thermal behavior
  - Conventional subarray scheme is not thermal-aware
- Proposed two subarraying schemes to optimize cache thermal behavior
  - Separated subarraying scheme
  - Way-interleaved subarraying scheme
- Even worsen thermal problem at deep sub-micron technologies demands a joint effort of thermal-aware subarraying schemes and leakage control

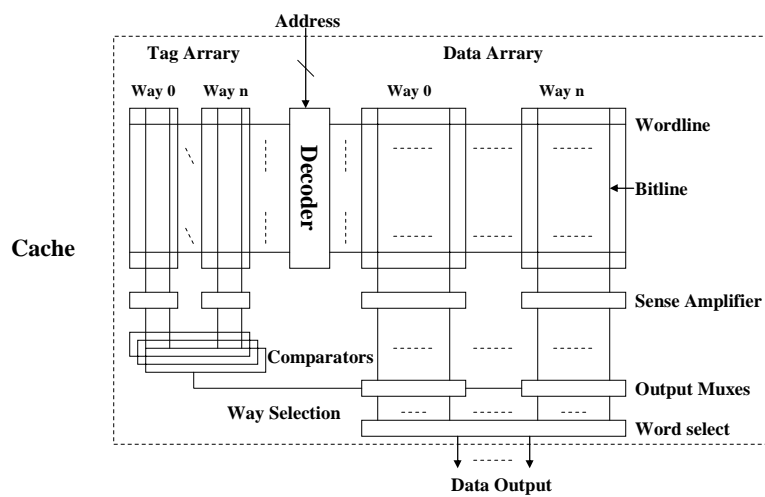
10/05/2005

ICCD 2005, San Jose

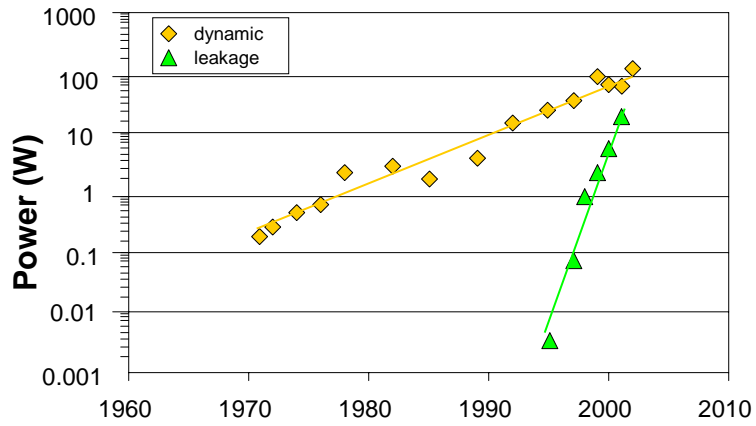
26

Thank you!

## Logical View of Set-Associative Caches



## Power: Dynamic vs. Leakage



Derived from Gordon Moore's Keynote Speech at ISSCC 2003.

10/05/2005

ICCD 2005, San Jose

29