Dependability, Power, and Performance Trade-off on a Multicore Processor

Toshinori Sato
Kyushu University
Toshimasa Funaki
Kyushu Institute of Technology

Outline

- What is soft errors?
- Multiple clustered core processor (MCCP)
- Power-performance trade-off
- Evaluation
- Conclusion

What is soft errors?

- Soft errors are random events which occur when various noise sources cause an incorrect event

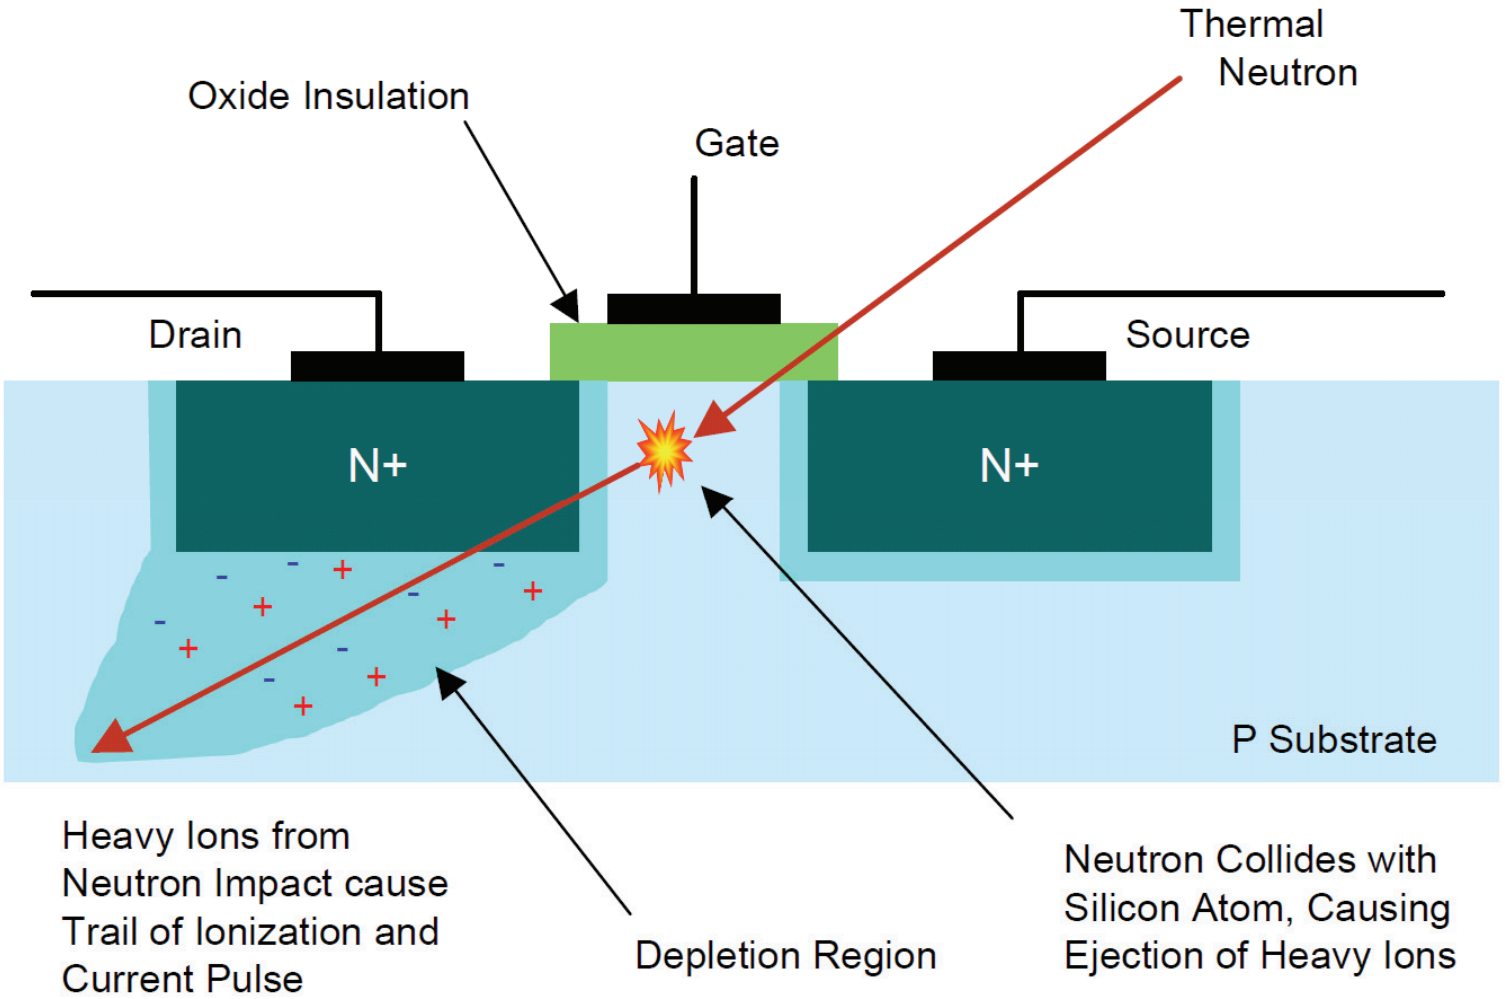

Source: Actel

Increasing soft error problem

- Advanced semiconductor technologies increase soft error rate

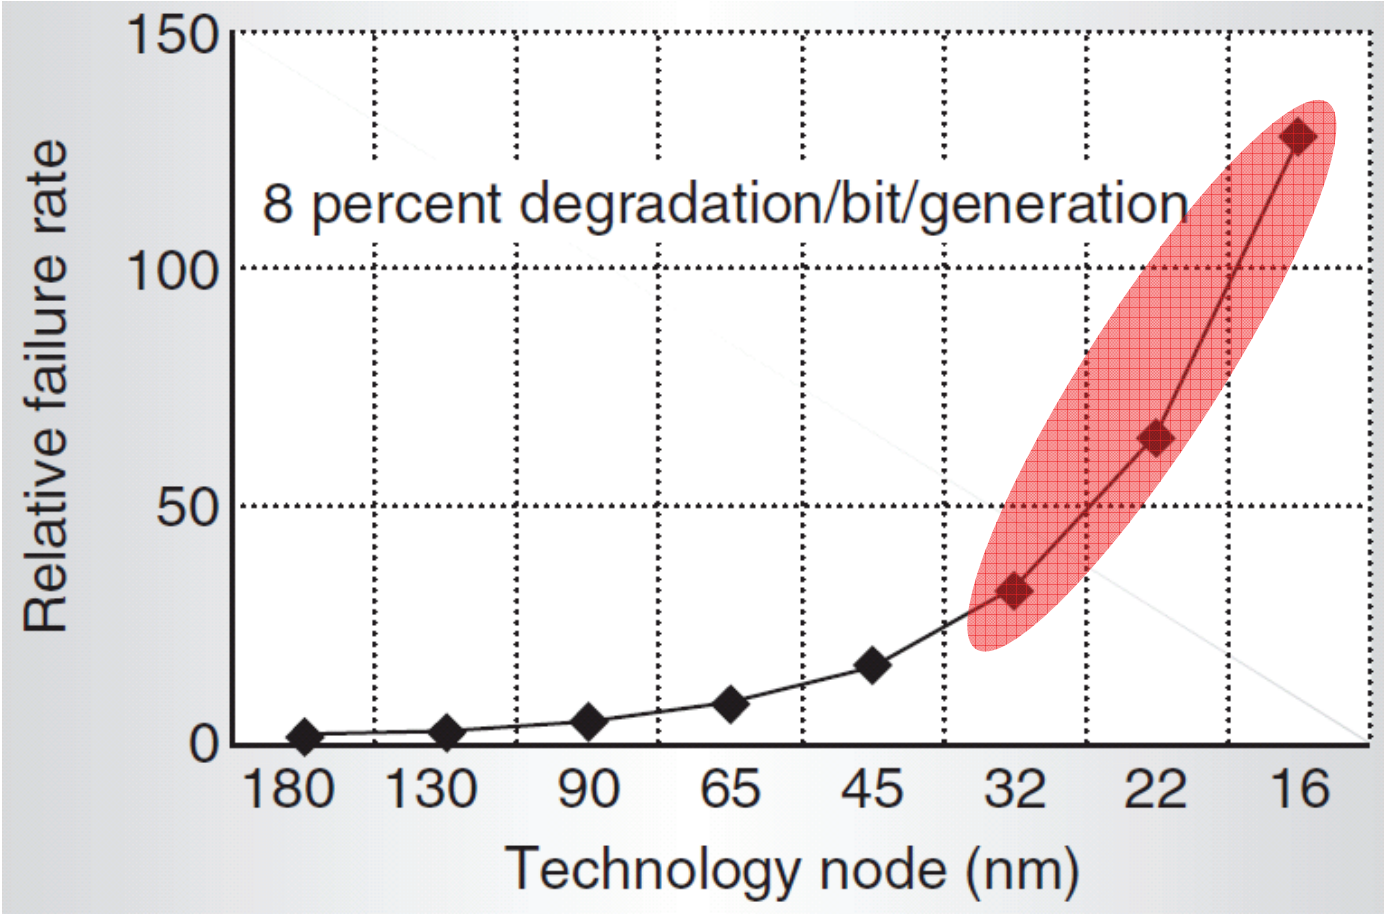

Source: Intel

Redundant threading (RT)

- Redundant execution is utilized
  - Space redundancy
  - Time redundancy
  - Focus on detecting soft errors (transient)
    - Error recovery is out of scope

Redundant threading (RT)

- Redundant execution is utilized
  - Space redundancy
  - Time redundancy
- Space redundancy
  - Thread-level redundancy
    - Popularity of multicore processors,
    - high performance
    - High power consumption
- Multiple clustered core processor
  - Considering power performance trade-off

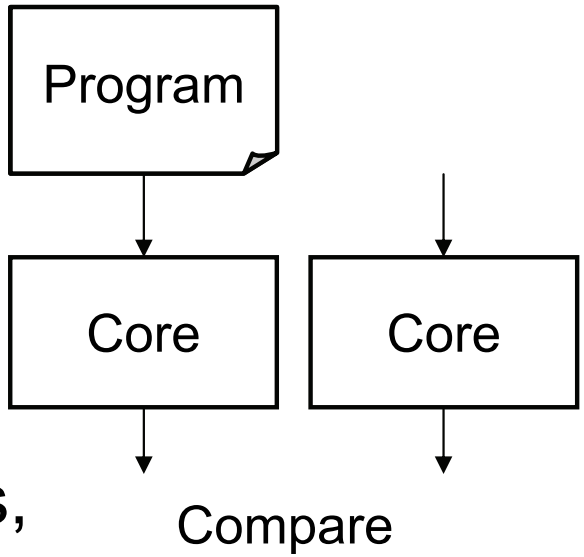

Multiple clustered core processor

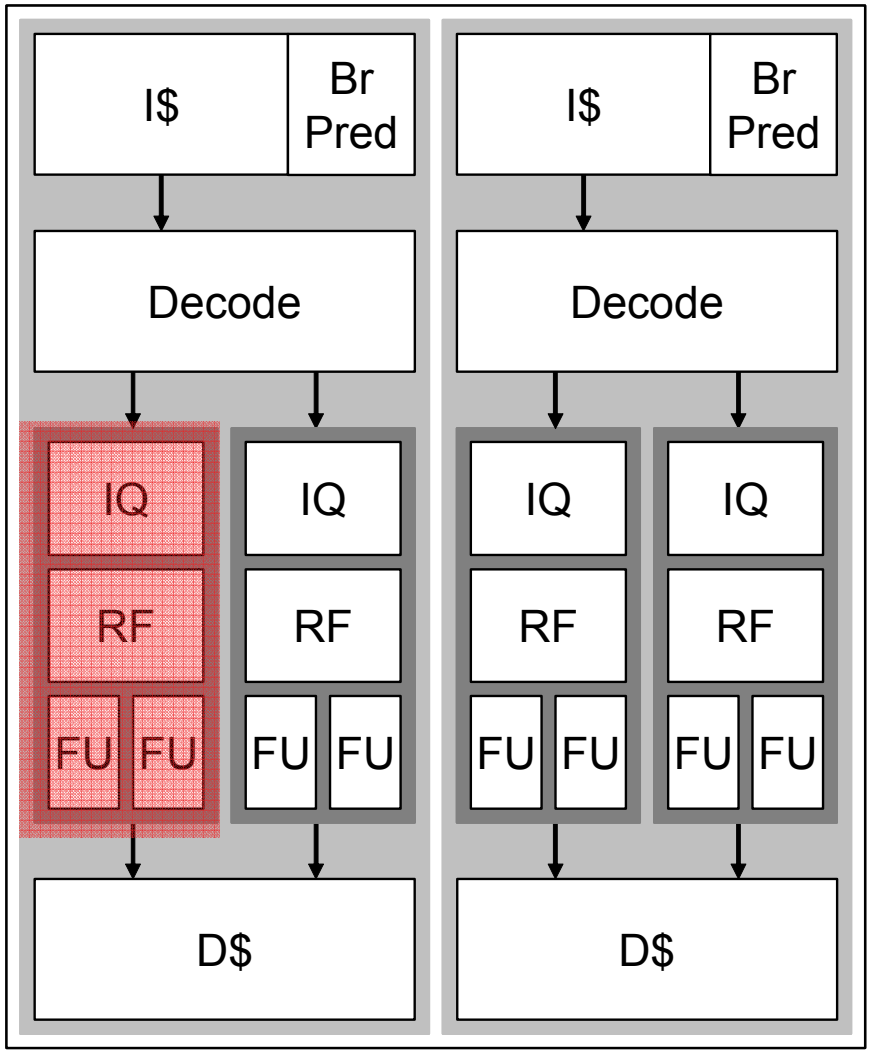

- Homogeneous
- Clustered cores
  - Datapath clusters
  - Configurable

Power-performance trade-off

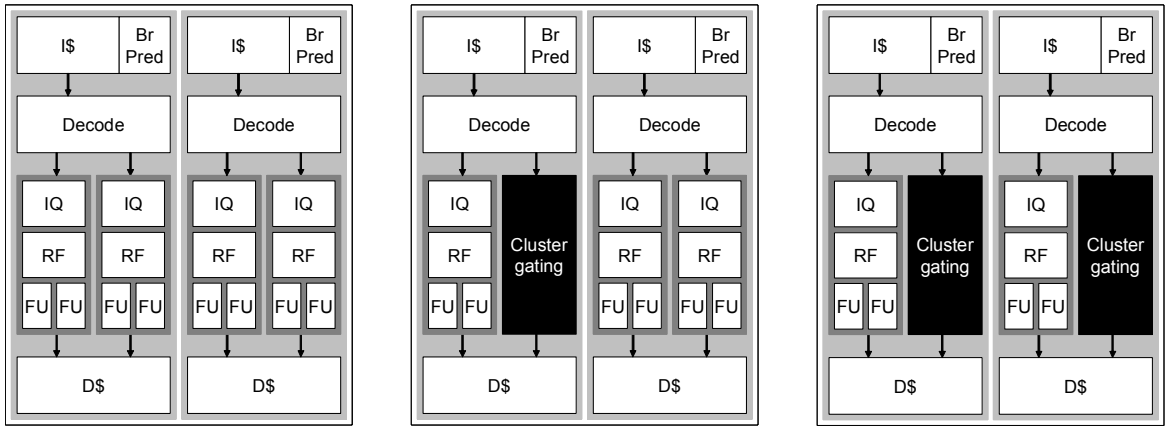

(a) Dual large core (b) Hetero core (c) Dual small core

- Cluster gating
  - Different scales and different configurations
- Consideration on trade-off between power and performance
  - Providing just required performance achieves power reduction

Pwr-perf on dependable proc.

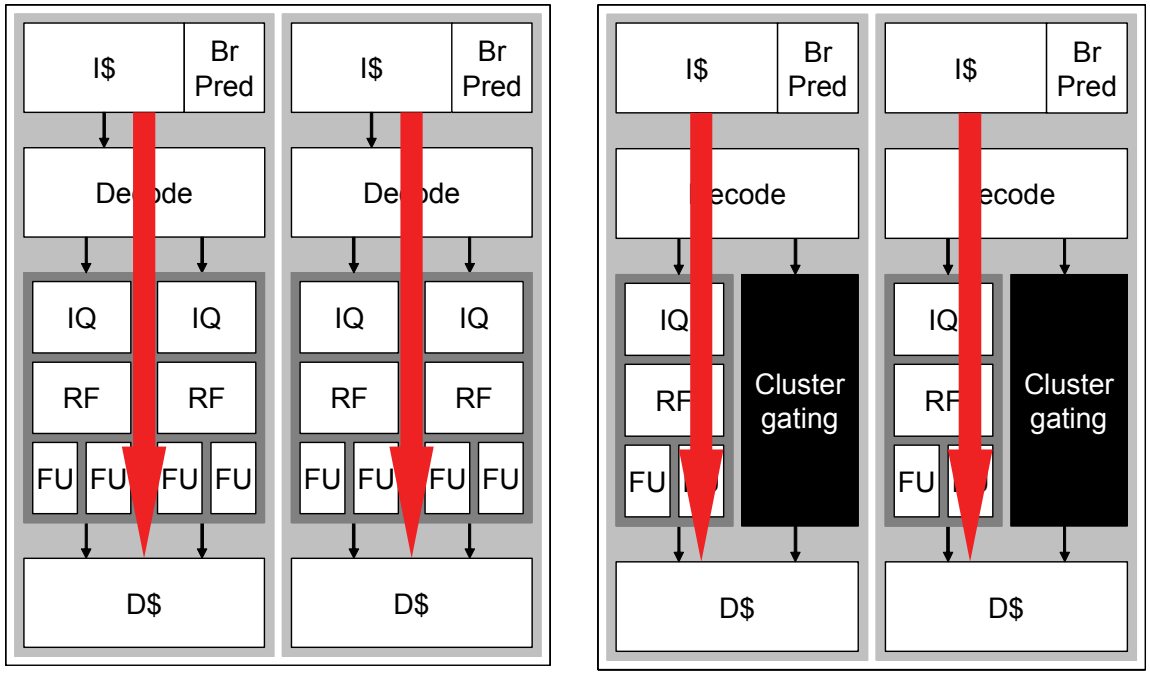

(a) High performance mode

(b) Moderate performance mode

- High performance mode
  - Power consumption is high
- Moderate performance mode
  - Power consumption is low

Pwr-perf on dependable proc.

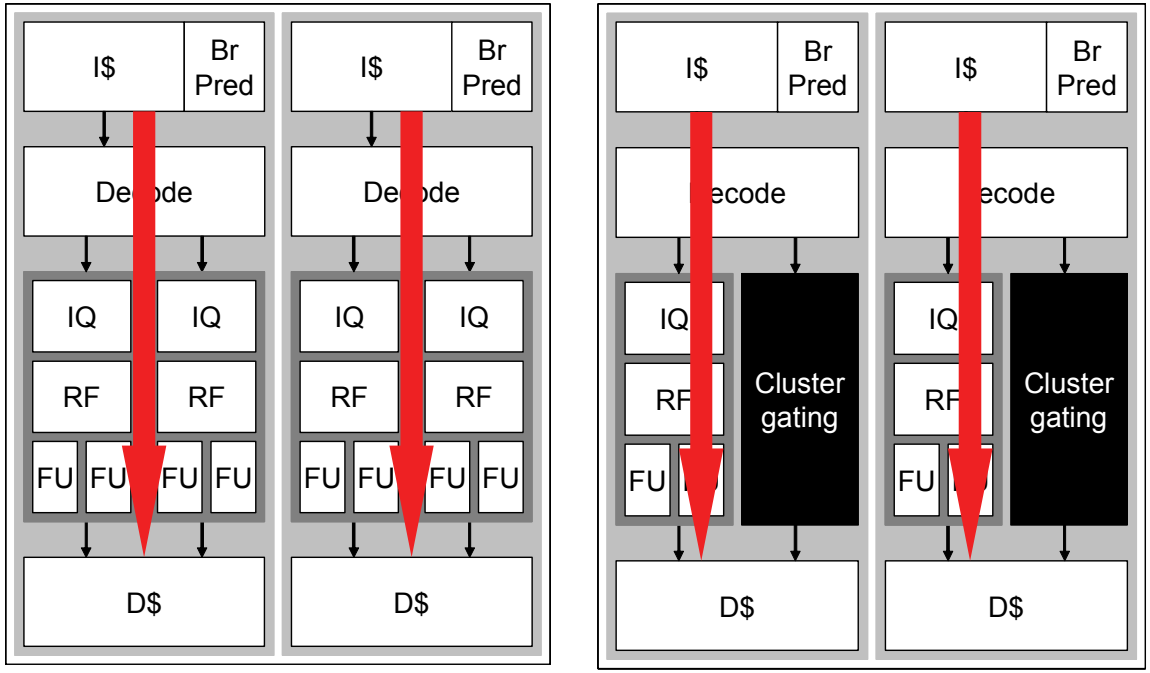

(a) High performance mode (b) Moderate performance mode

- High performance mode
- Low performance mode
  - Power consumption is low

How do we determine the mode?

How do we determine mode?

- Annotation-based
  - A programmer marks how important every instruction block is and tells it to HW using annotations
- OS-based
  - OS marks the importance of every instruction block using some metric; ex. deadline time
- HW-based
  - HW transparently evaluates the importance of every instruction block

Issue IPC variation

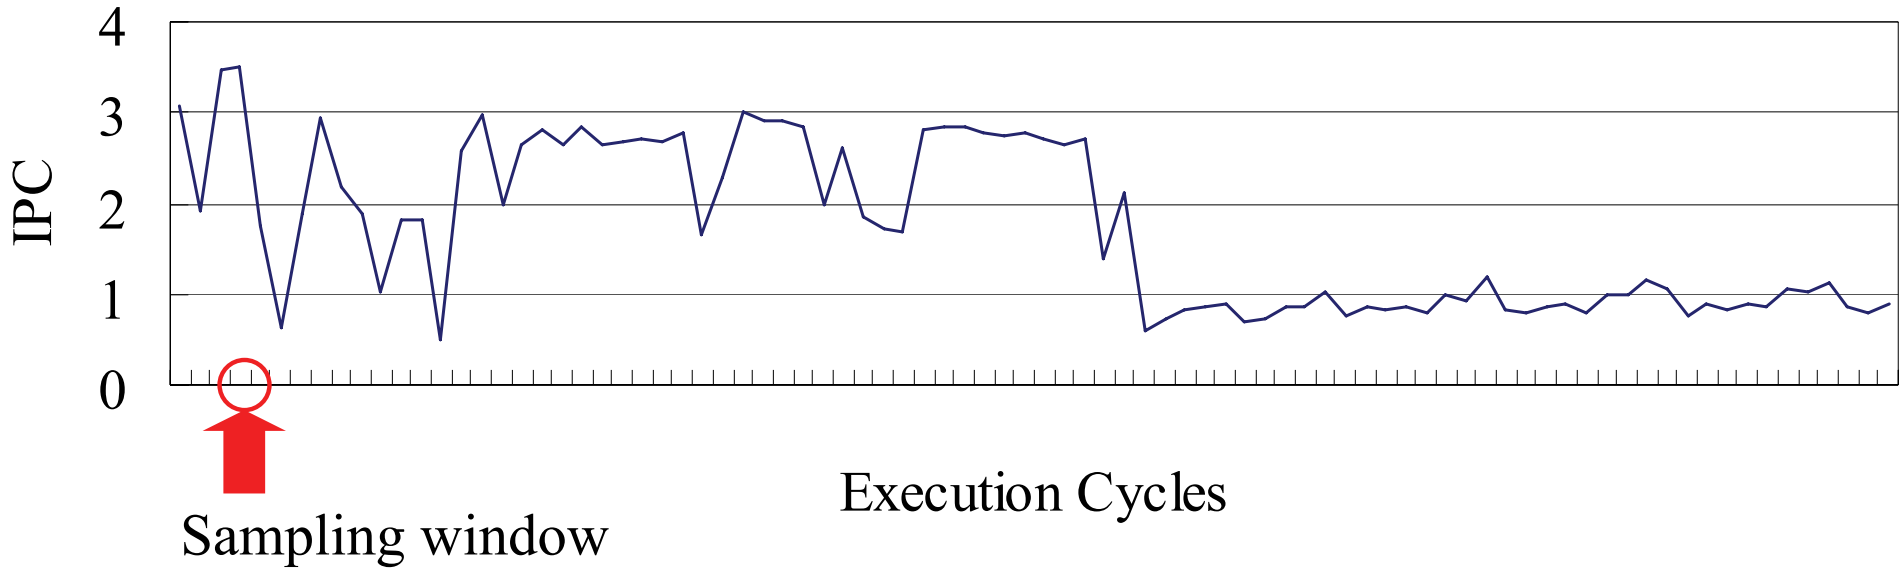

- Issue IPC changes by more than a factor of 2
- The variations can be used to determine the mode

IPC-directed mode selection

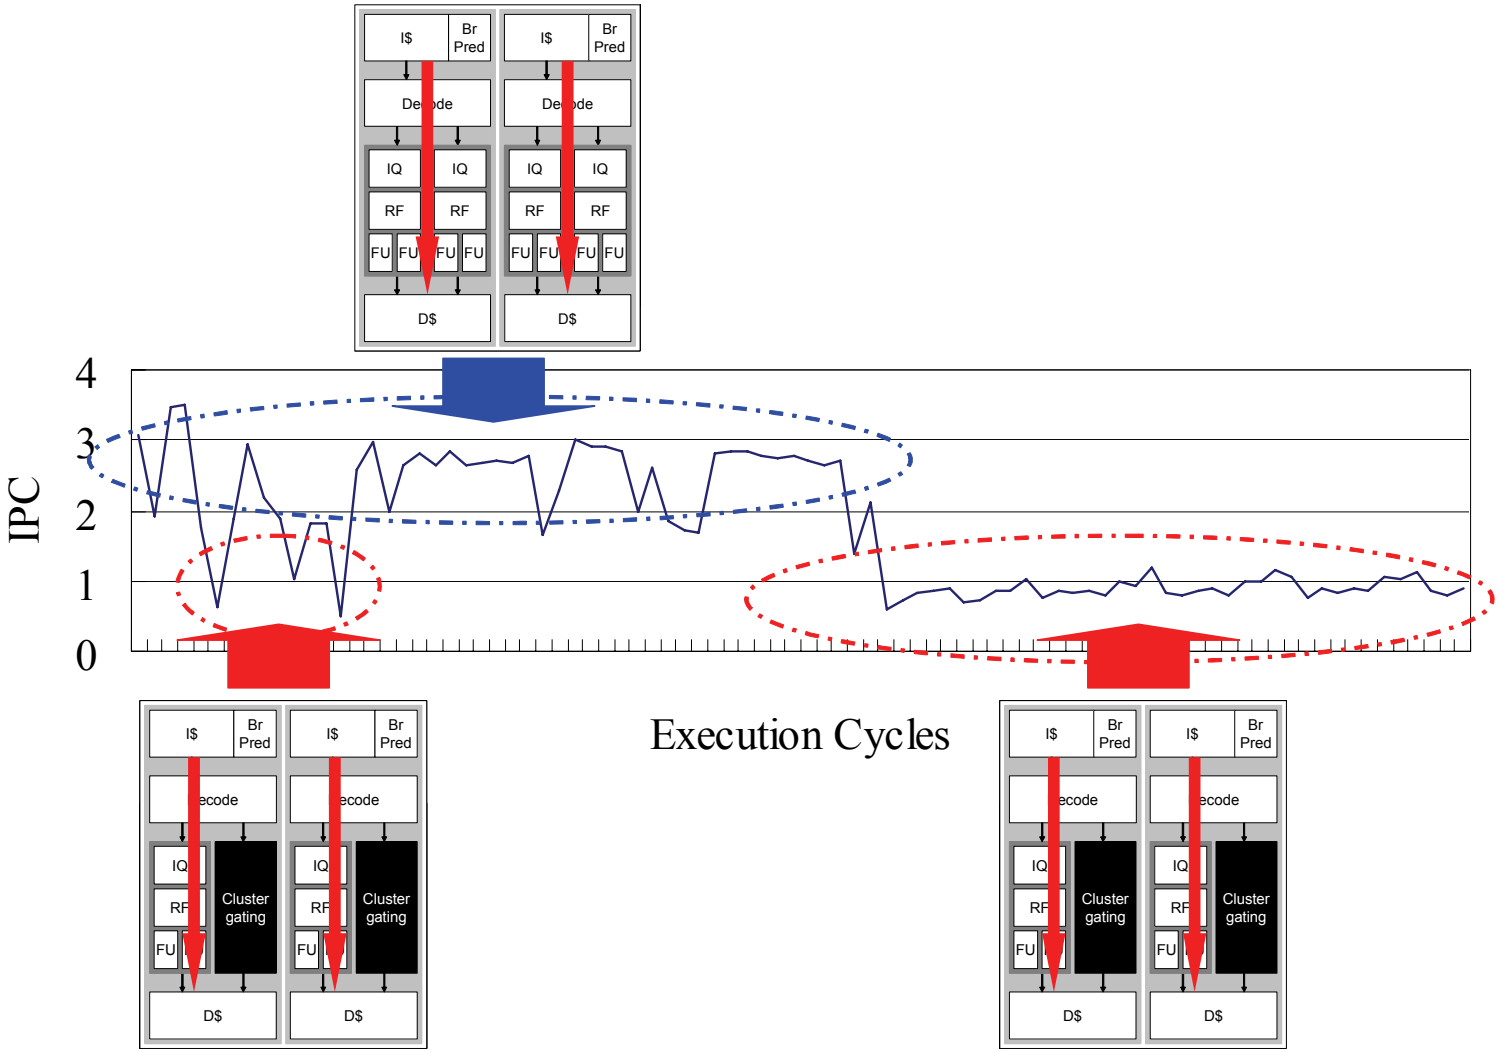

Future IPC prediction

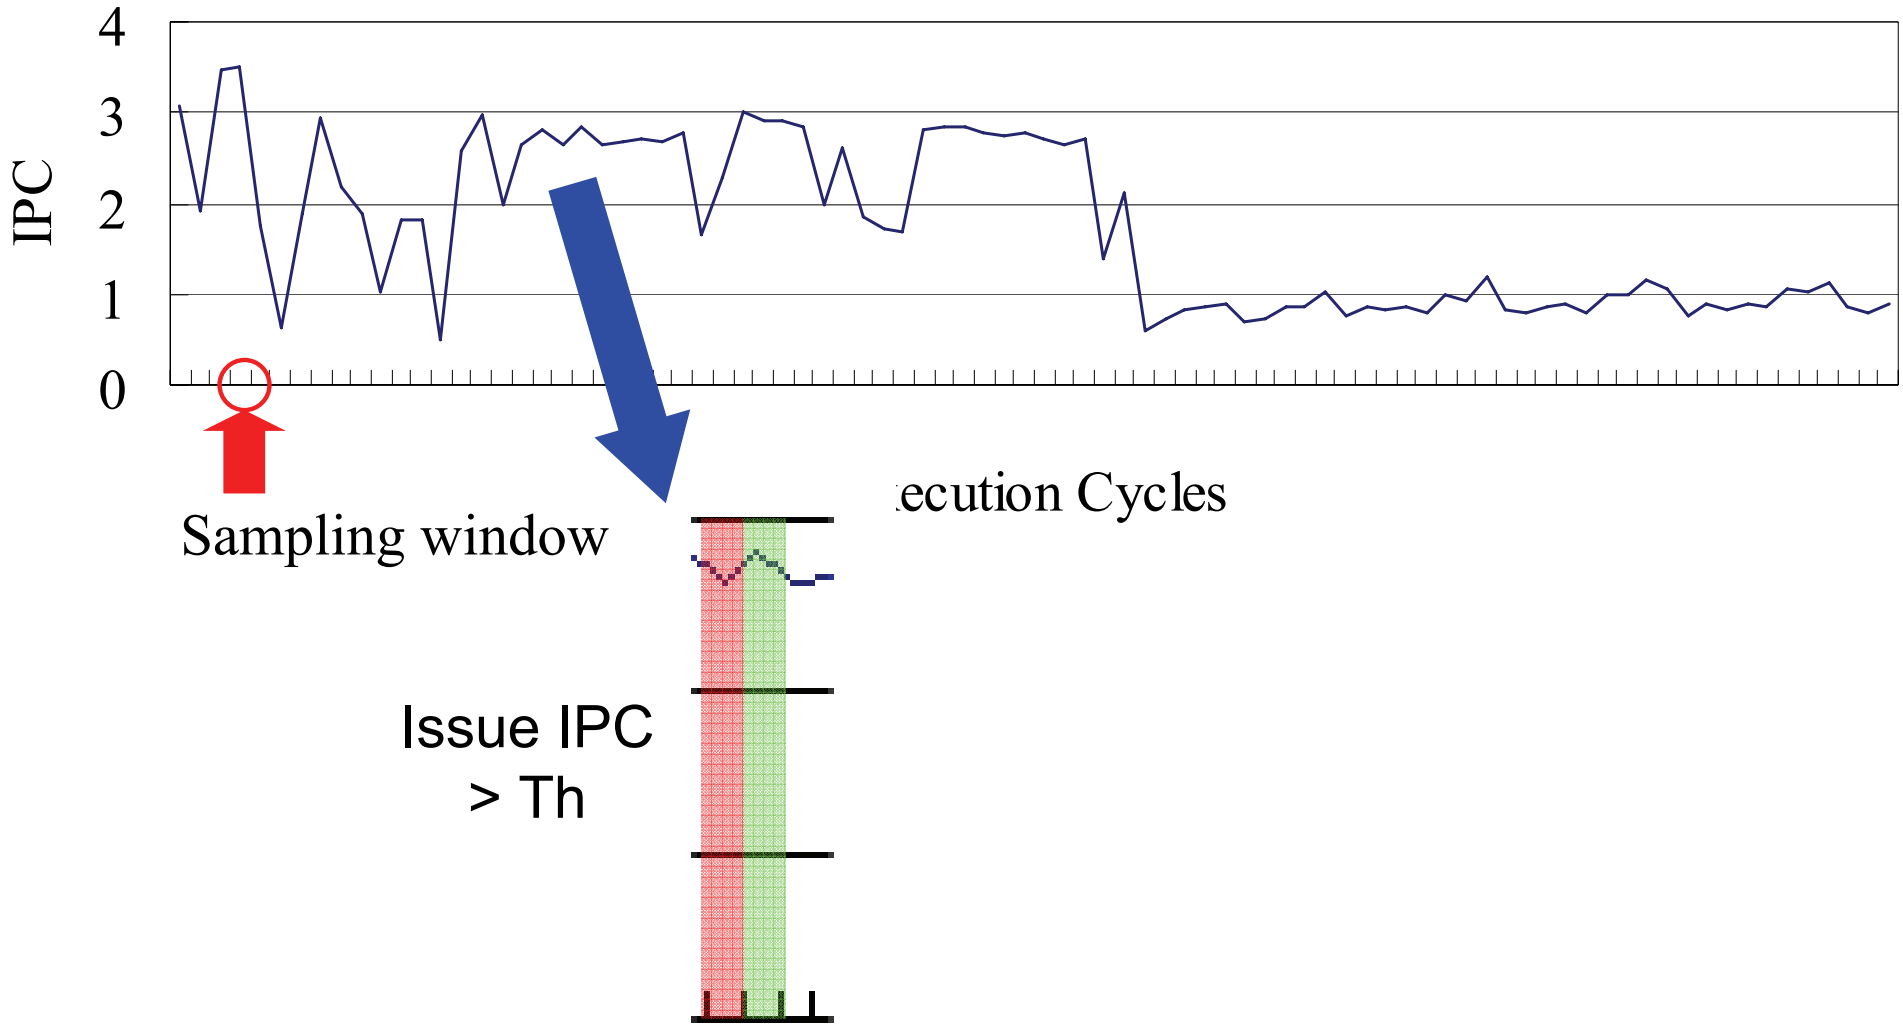

Future IPC prediction

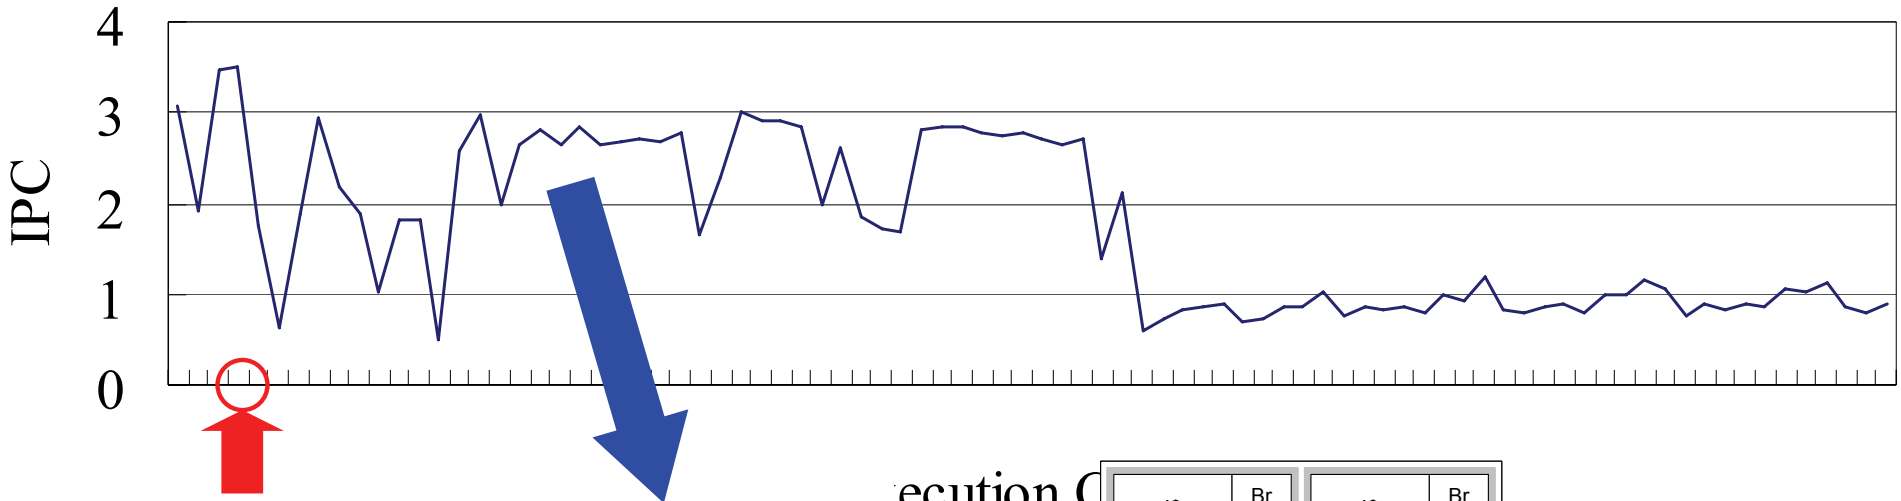

Sampling window

Issue IPC
> Th

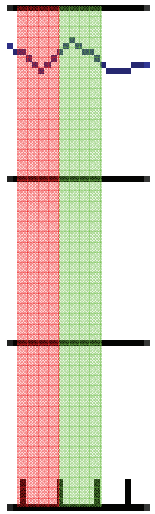

ecution C

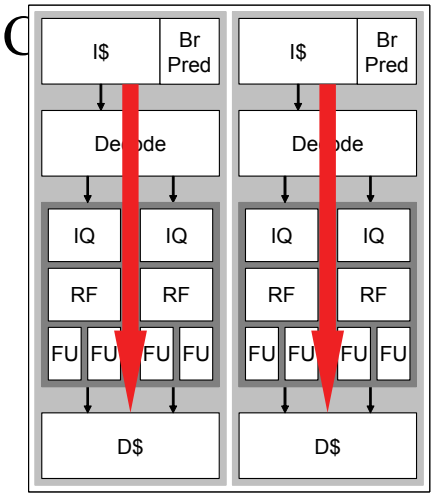

High performance
mode is selected

Methodology

- SimpleScalar/PISA
- Dual dual-cluster core MCCP
- Mode switch threshold
  - 2.0: High -> Moderate
  - 1.6: Moderate -> High
- Benchmark
  - SPEC2K
  - MediaBench
- Sampling window
  - 100, 1,000, 10,000

Fetch width	8 instructions
L1 instruction cache	16K, 2-way, 1 cycle
Branch predictor	1K-gshare + 512-BTB
Dispatch width	4 instructions
Instruction window size	16 entries / cluster
Issue width	2 instructions / cluster
Commit width	4 instructions / cluster
Integer ALUs	2 units / cluster
Integer multiplires	2 units / cluster
Floating ALUs	2 units / cluster
Floating multiplires	2 units / cluster
L1 data cache ports	1 port / cluster
L1 data cache	16K, 2-way, 1 cycle
Unified L2 cache	512K, 2-way, 10 cycles
Memory	Infinite, 100 cycles

Upset rate estimation

- Upset rate is estimated as a product of area and soft error rate per bit
- Shared components are not included for estimation
- High : Moderate = 100 : 72

16K L1 data cache	2.6
16K L1 instruction cache	2.6
TLB	4.4
Fetch unit	1.3
Branch predictor	3.2
Decoder	1.7
OOO execution unit	10.1 / cluster
Register files	2.9 / cluster
Functional units	6.5 / cluster
Misc	2.4
Routing	26.4
512 L2 cache	110.0
Misc	6.1
Coherence unit	6.3
I/O	13.7

Power assumption

- Power is estimated using 21264 power consumption
- High : Moderate = 100 : 81

Global Clock Network	32%
Instruction Issue Units	18%
Caches	15%
Floating Execution Units	10%
Integer Execution Units	10%
Memory Management Unit	8%
I/O	5%
Miscellaneous Logic	2%

Table 2: Alpha 21264 Power Components

M Gowan, Power considerations in the design of the Alpha 21264 microprocessor, DAC (1998)

Metric

- To evaluate the trade-off between power, performance, and dependability, a new metric is required
- EDP (energy-delay product) is a popular metric for power-performance trade-off.
- Upset rate is a popular metric for dependability
- We propose to product EDP and upset rate
  - EDUP (energy, delay, and upset-rate product)

Breakdown of dependability mode

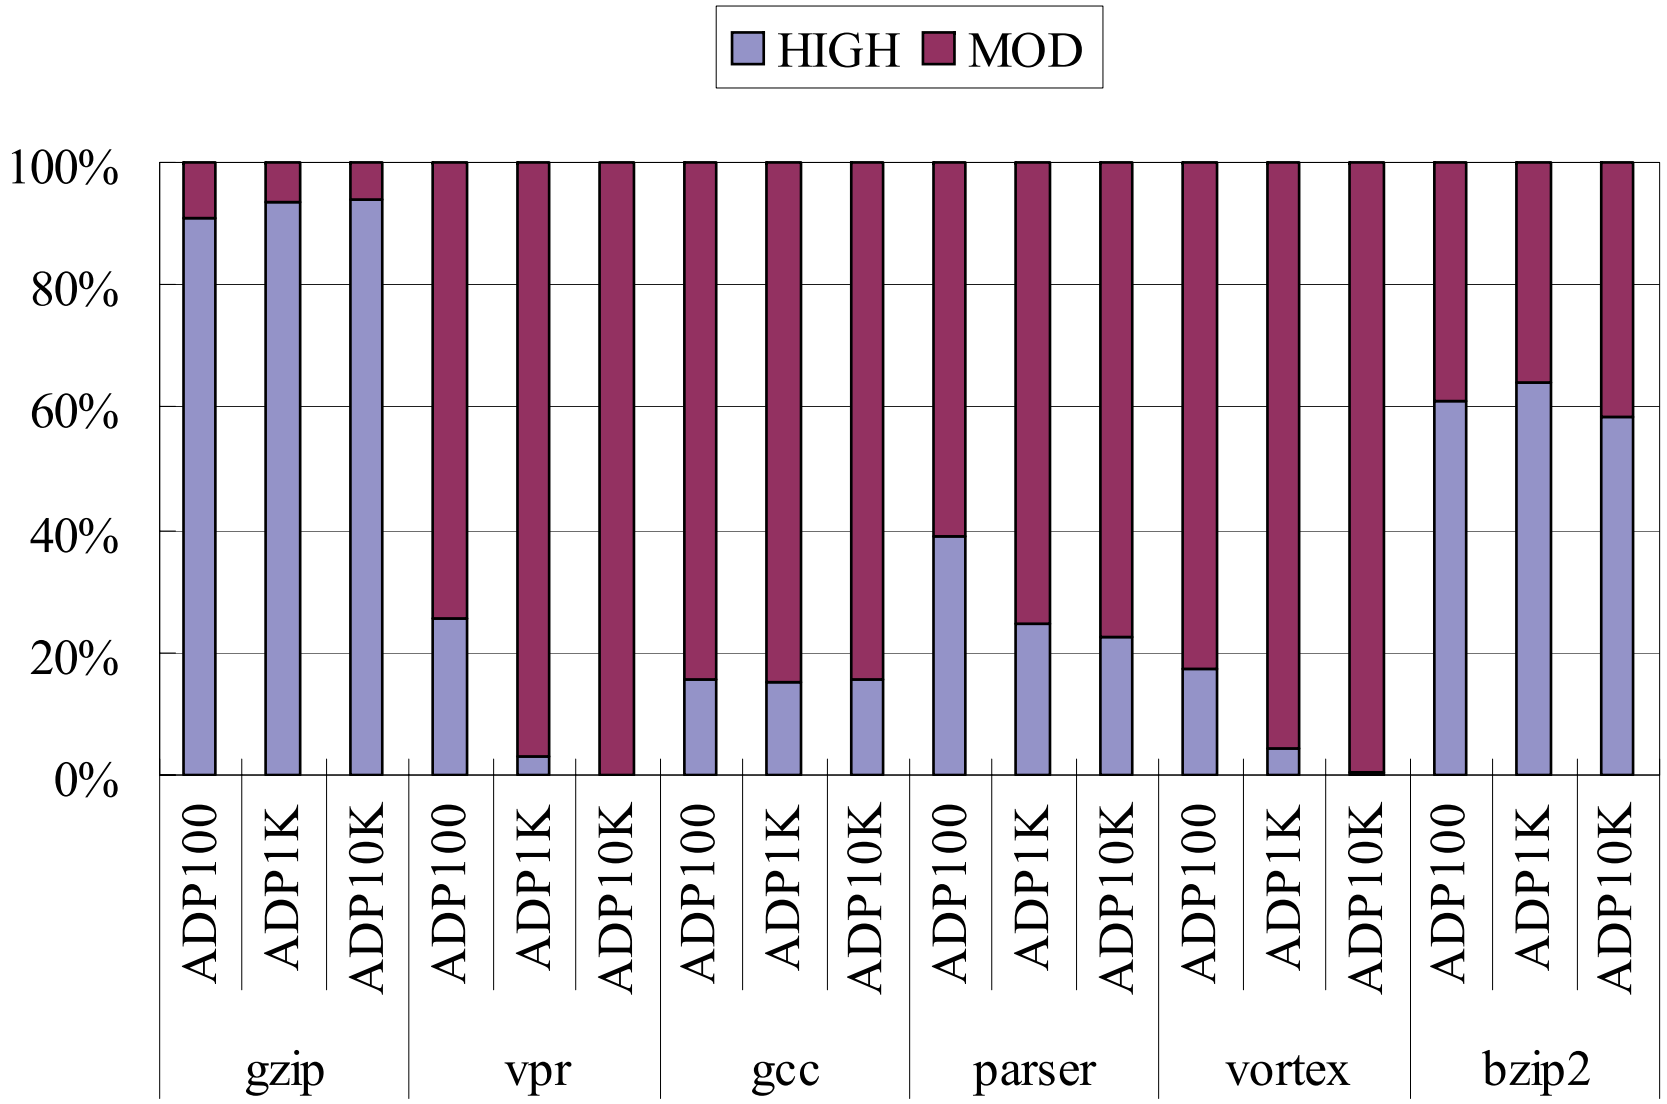

Breakdown of dependability mode

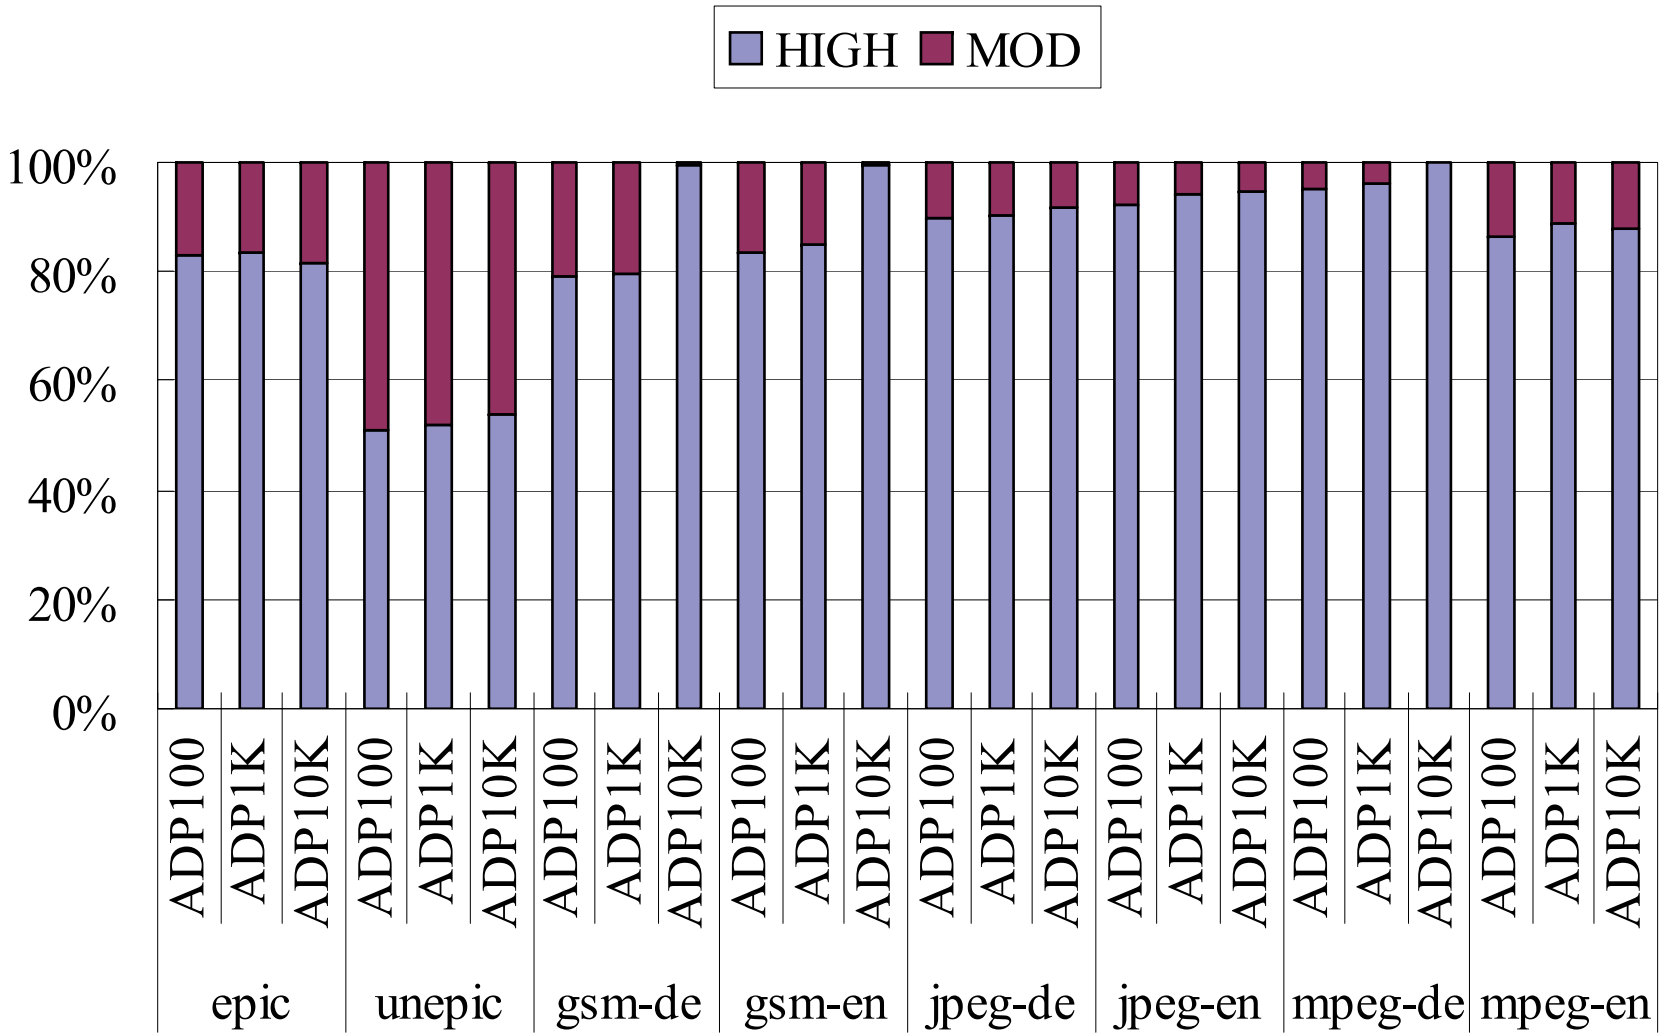

Relative IPC and EDUP

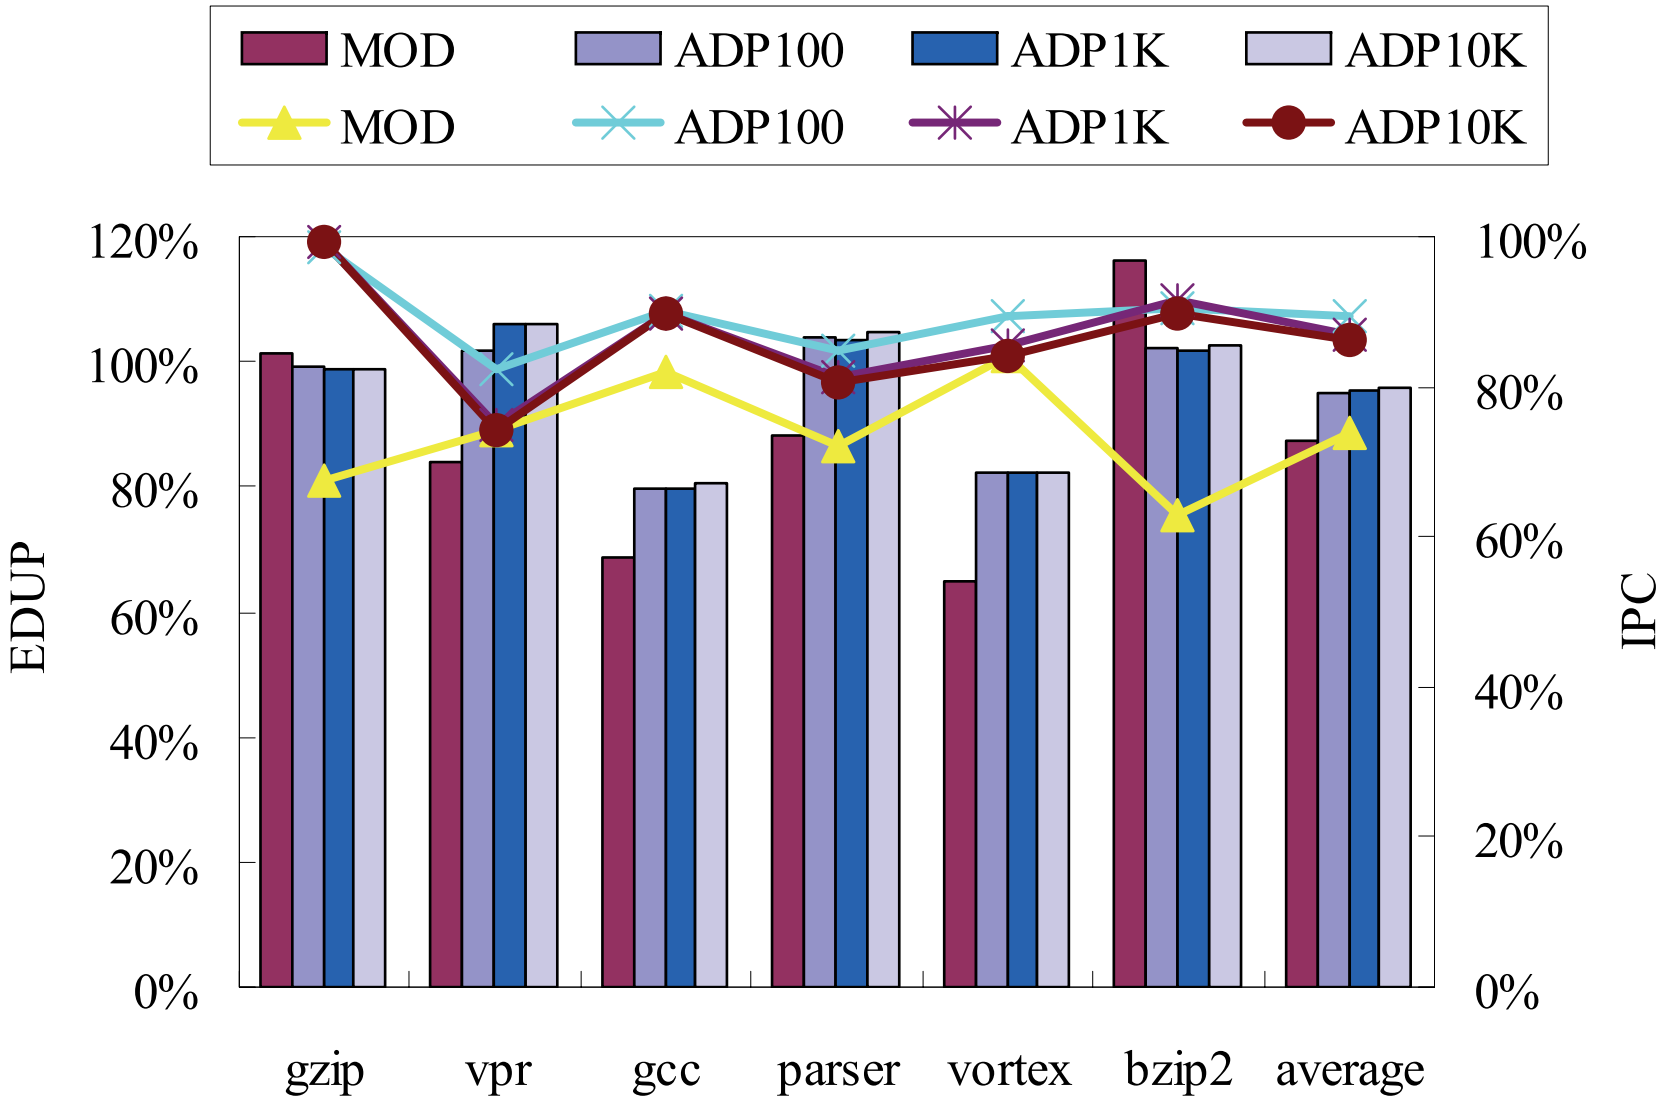

Relative IPC and EDUP

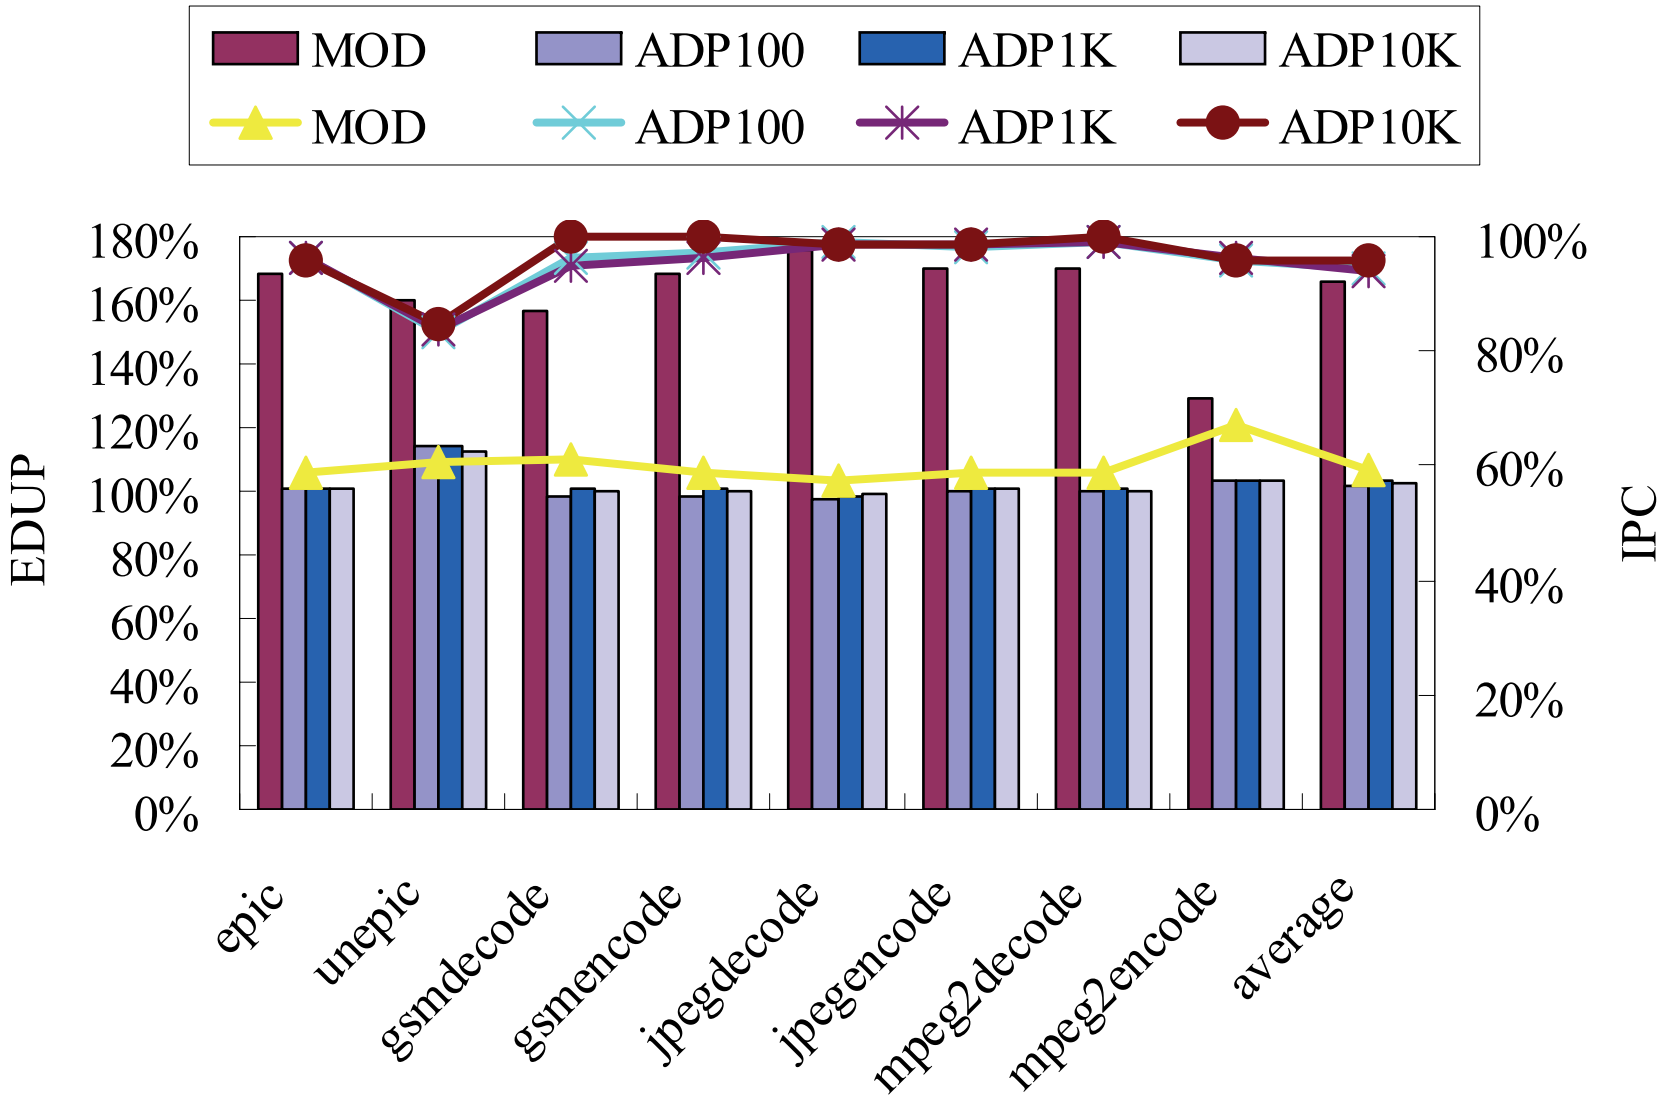

Conclusions

- RT is widely utilized for soft error tolerance, but it consumes much power
- We proposed an adaptable RT, named MCCP
  - It has high- and moderate-performance modes for considering dependability, power, and performance trade-off
- We also proposed EDUP as a metric for evaluating the trade-off
- Detailed simulation showed MCCP improved EDUP by up to 21%

Dependability, Power, and Performance Trade-off on a Multicore Processor

Toshinori Sato
Kyushu University
Toshimasa Funaki
Kyushu Institute of Technology