



Designing energy-efficient microprocessor: How to fight process variations

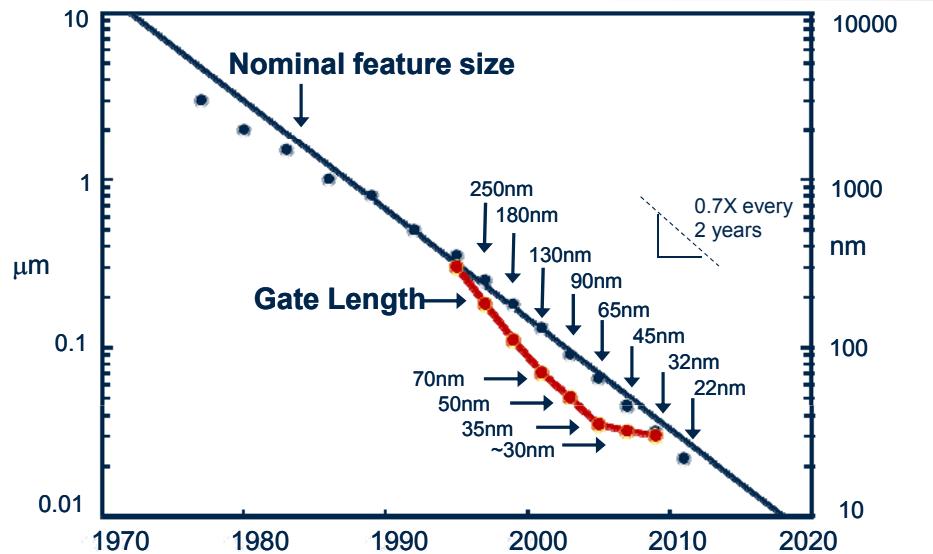
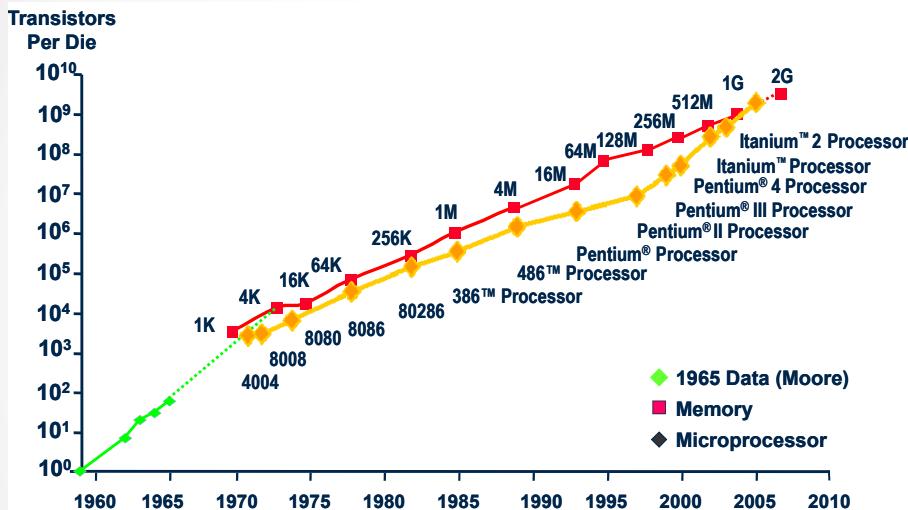
Ruzica Jevtic
Berkeley Wireless Research Center



POLITÉCNICA



Moore's law

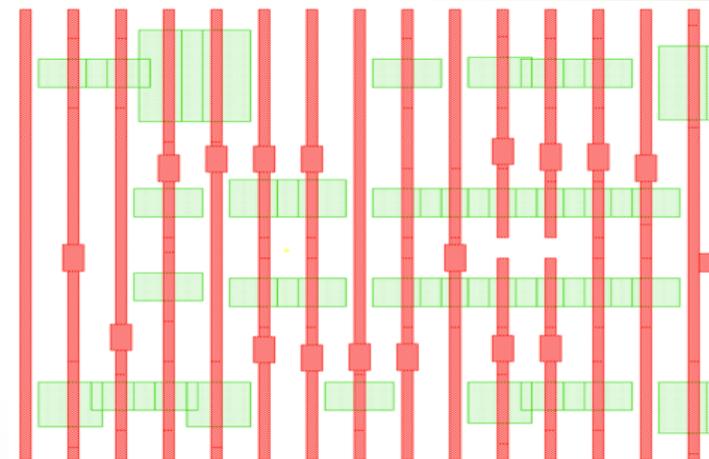
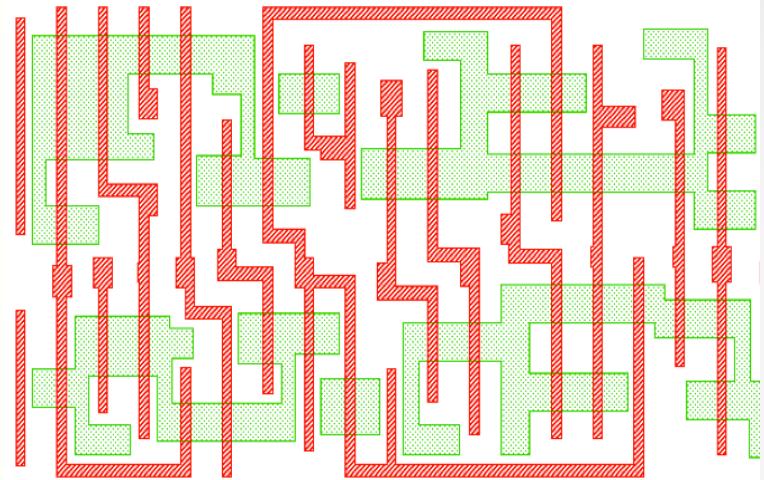


Scaling consequences:

- Process variations
- Power has become critical!

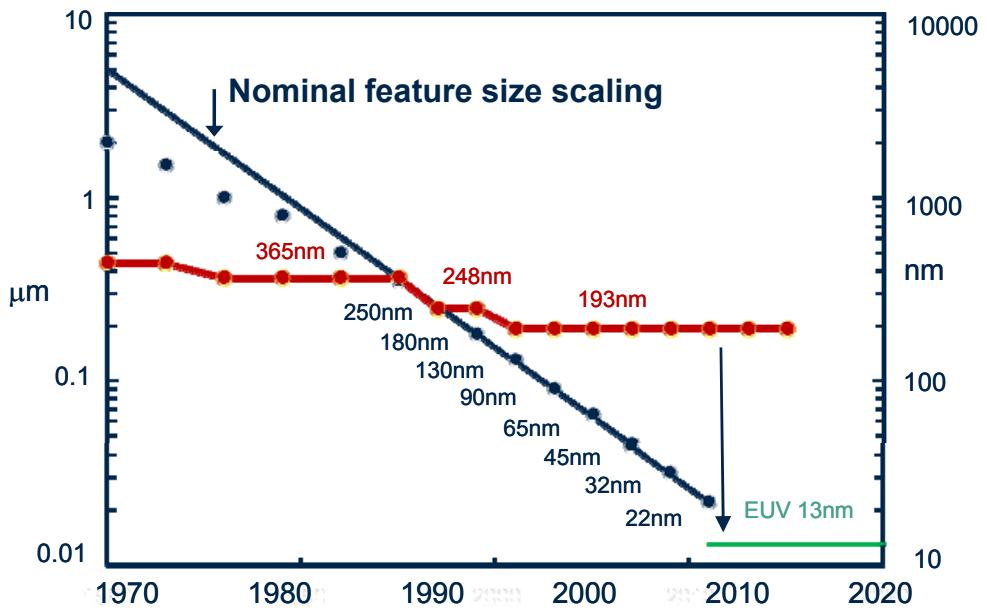
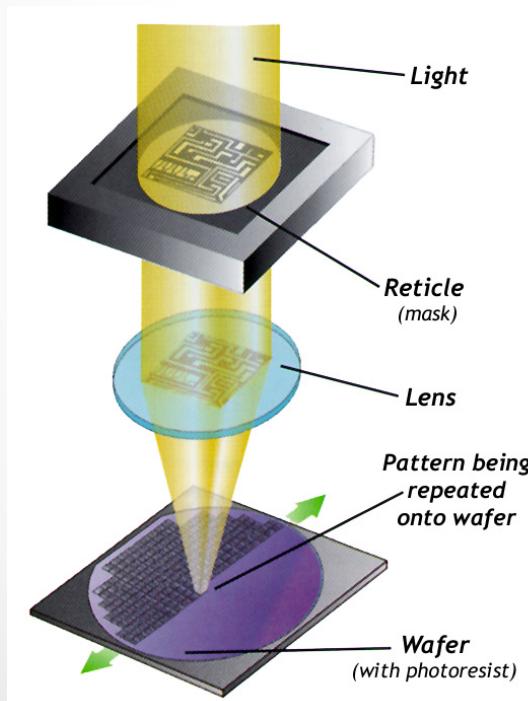
Scaling consequences

- **Fabrication:**
 - Litography with liquids
 - Use of diffraction...
- **Design rules:**
 - No corners
 - One direction lines...
- **Additional steps:**
 - Optical Proximity Check
- **Consequences:**
 - Random dopant fluctuations
 - Hot-carrier trapping...

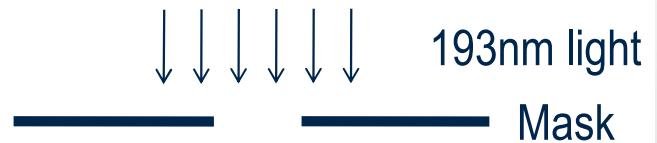


Litography

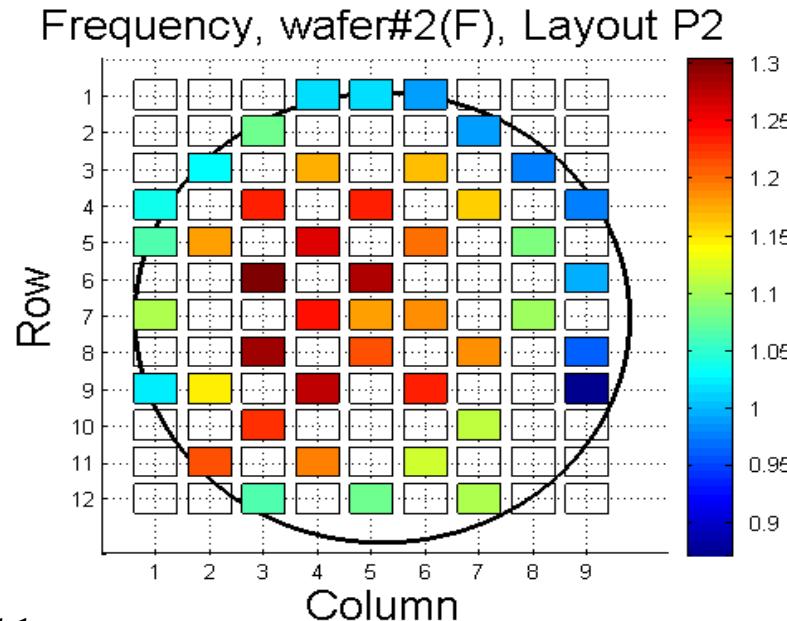
- Litography issues:
193nm wavelength for
lines as small as 30nm!



Light projected through a gap



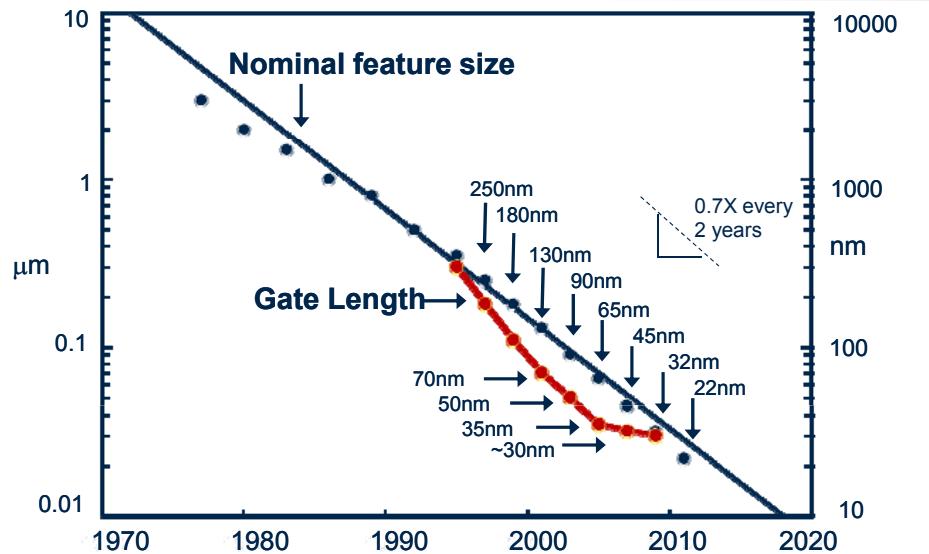
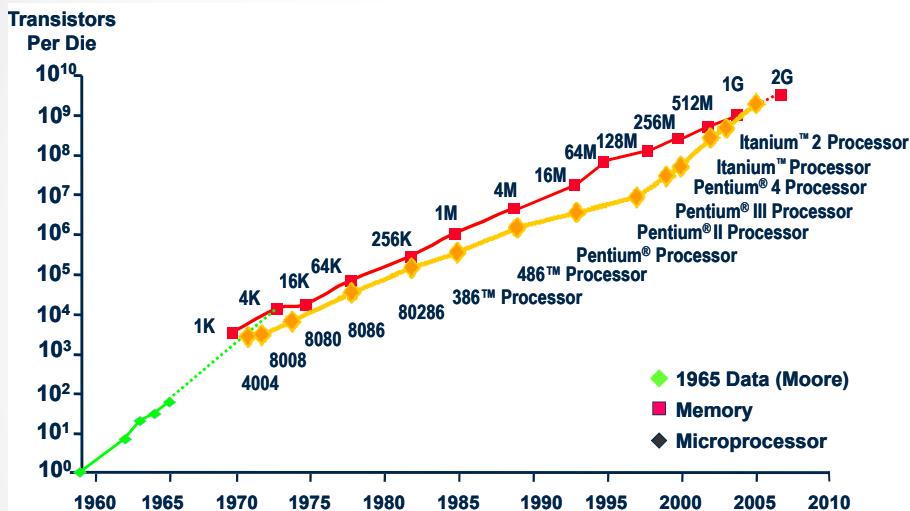
Process variations



B. Nikolic, TCAS-I, 2011

- Process corners: typical, fast and slow
- All corners coexist on the same wafer
- If we design for the worst case, it is too pessimistic!
- Need to know after fab chip features → observation circuits

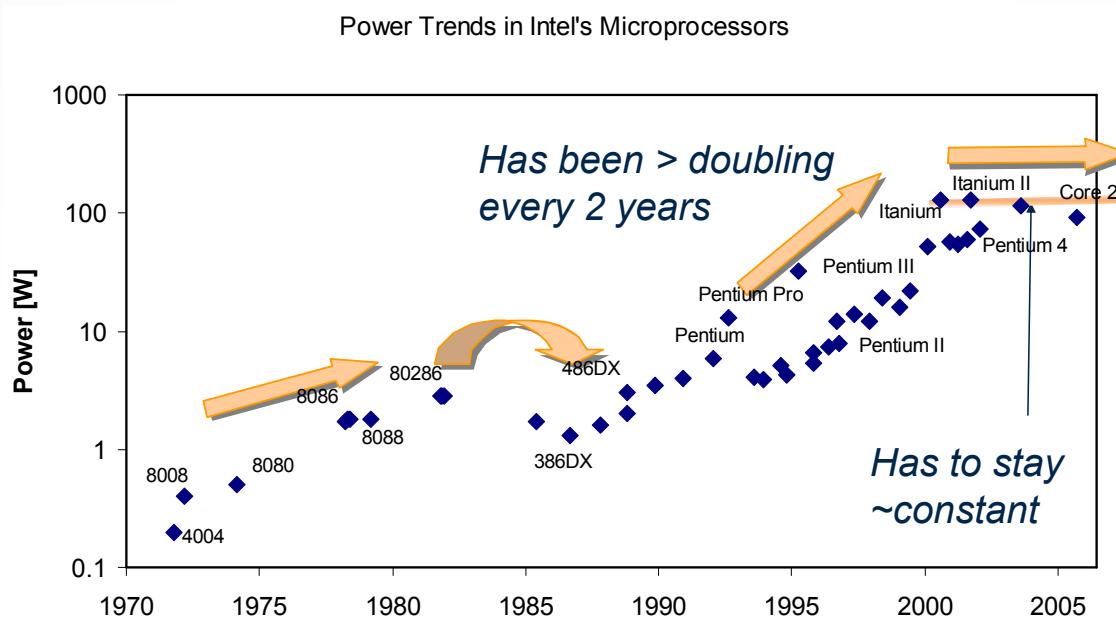
Moore's law



Scaling consequences:

- Process variations
- Power has become critical!

Scaling issues: Power



- Operating voltage scaled to avoid device breakdown
- Not scaled enough to keep the performance boosting up
- Power has become a critical design constraint

Power issues

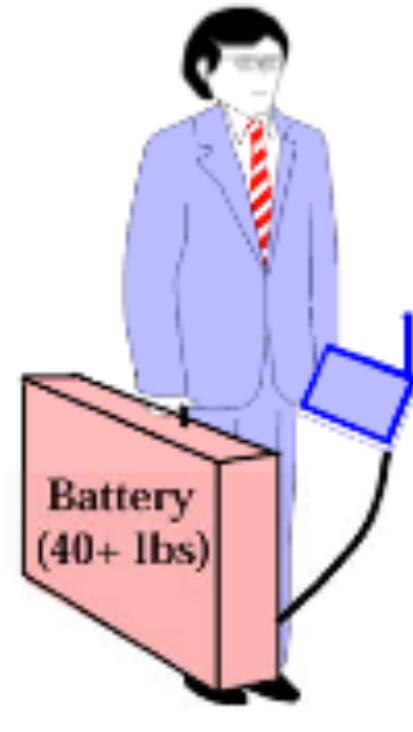
- Power became an issue before variations:
 - battery life for mobile devices
 - power limits performance!
 - performance is what sells the product!

- Power classification:

- Dynamic power (60%)

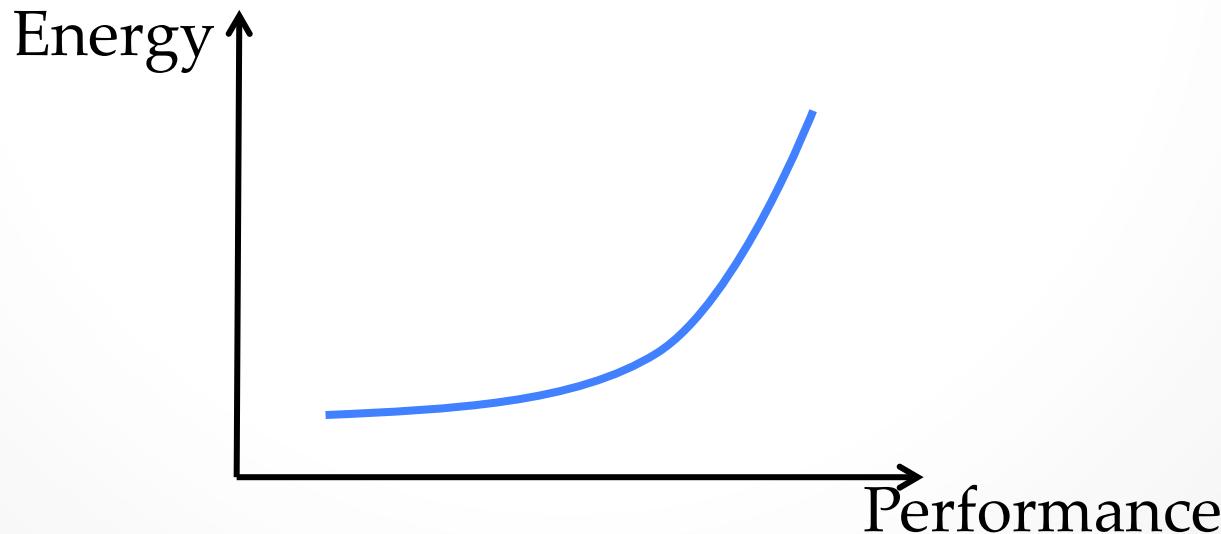
$$P = SW \cdot f \cdot V_{dd}^2 \cdot C$$

- Static power (30%) $\sim \exp(V_{dd})$



Power-performance trade-offs

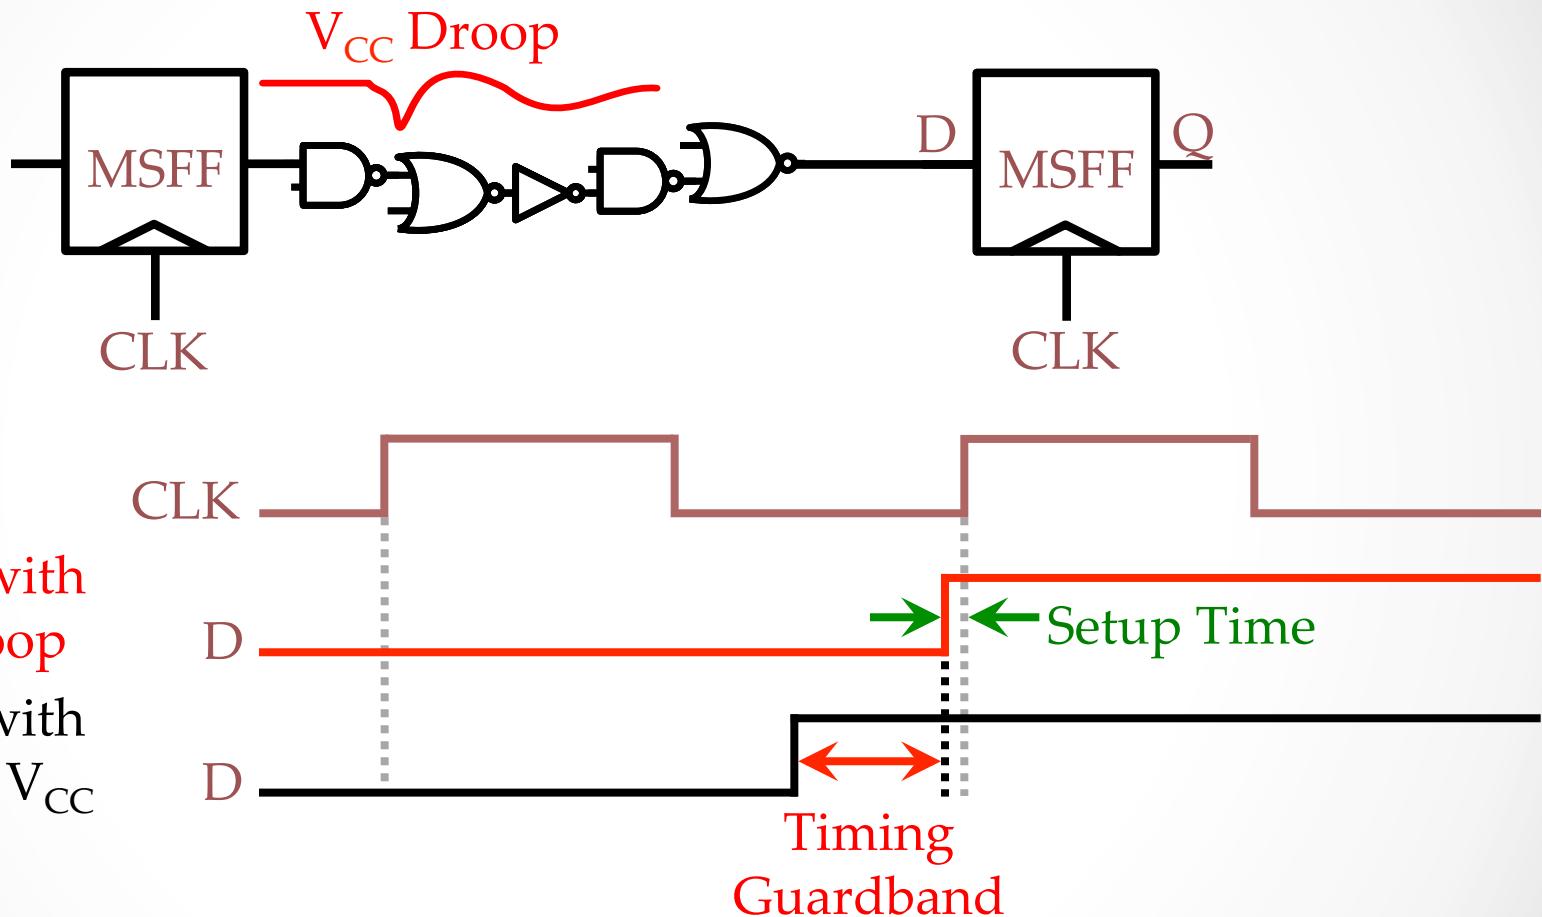
- The best way to reduce power (both leakage and active) is to reduce the power supply
- How to maintain throughput under reduced supply?
- Introducing more parallelism/pipelining
- Dynamic voltage scaling with variable throughput



Overview

- Error detection and correction circuits
- Dynamic voltage and frequency scaling
- Control voltage through DC-DC converters
- Conclusions

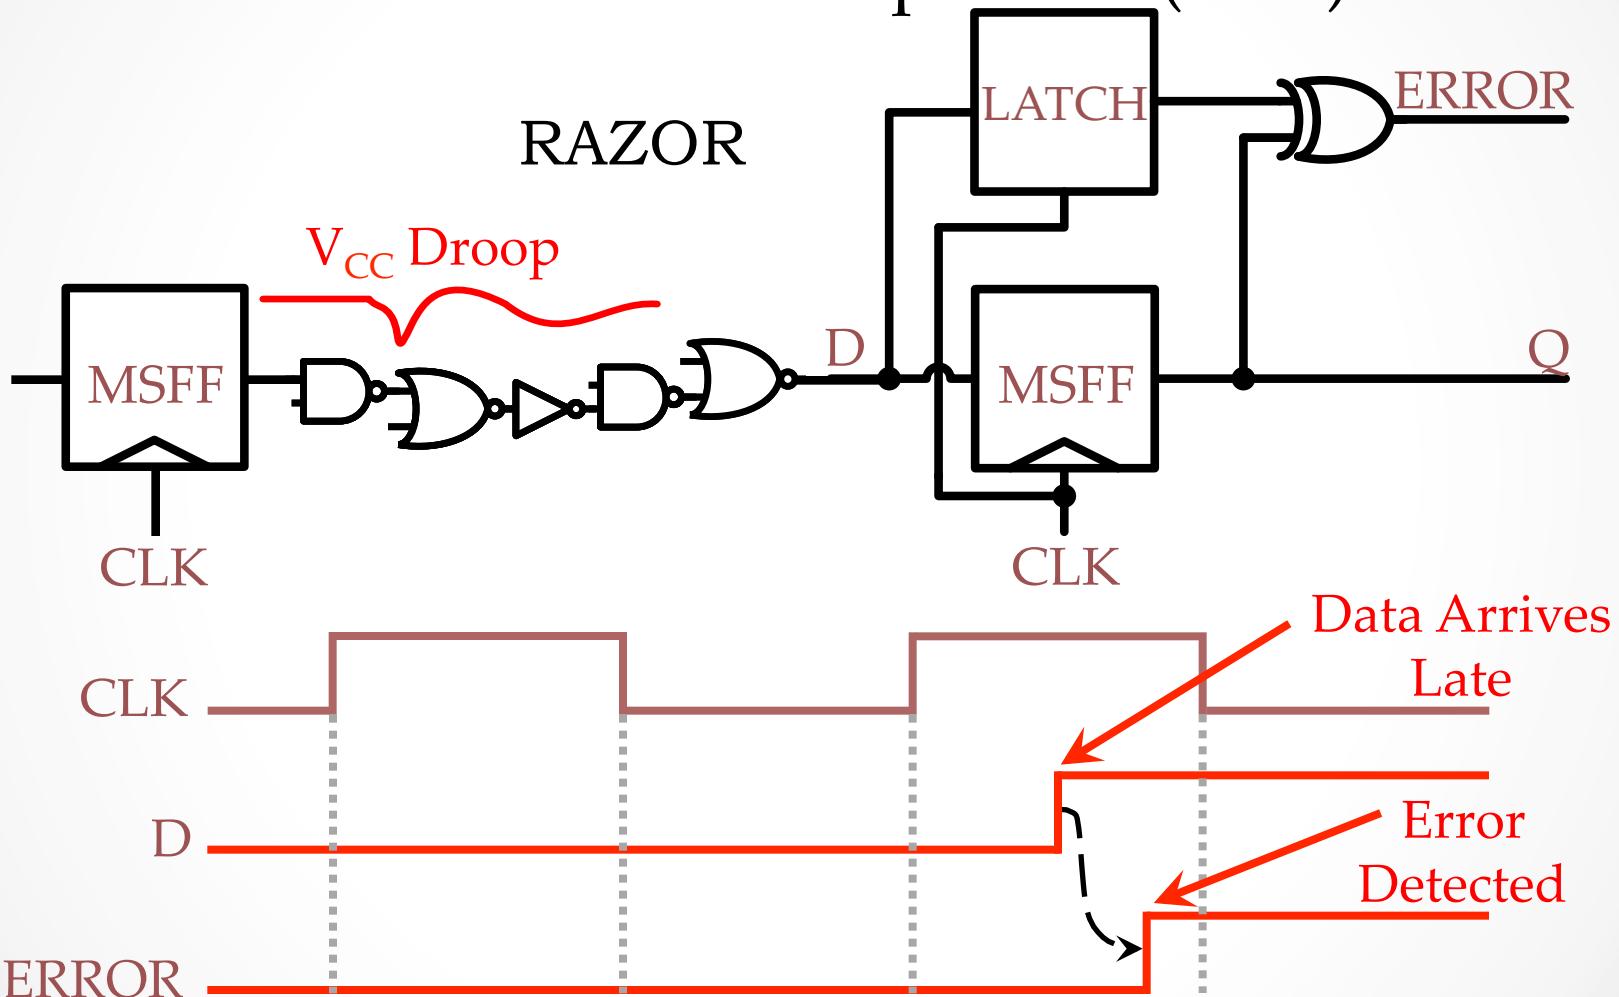
Impact of Dynamic Variations



- Guardbands required to ensure correct operation within the presence of dynamic variations

Timing-Error Detection

Error-Detection Sequential (EDS)

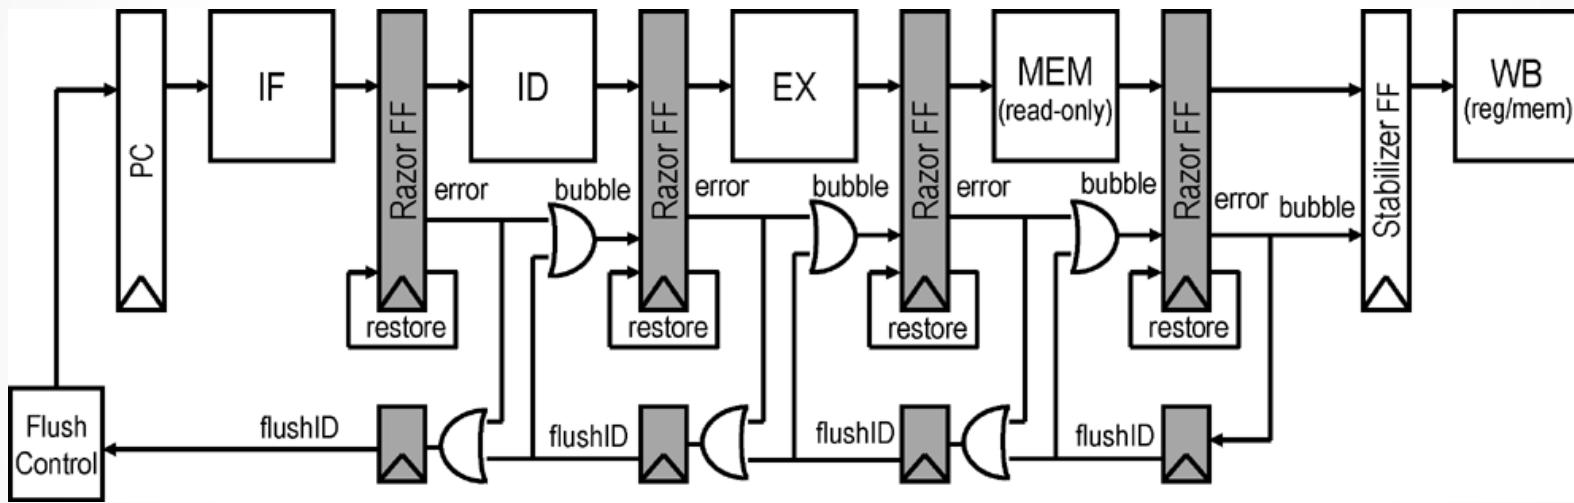


[1] P. Franco, et al., *VLSI Test Symp.*, 1994.

[2] M. Nicolaidis, *VLSI Test Symp.*, 1999.

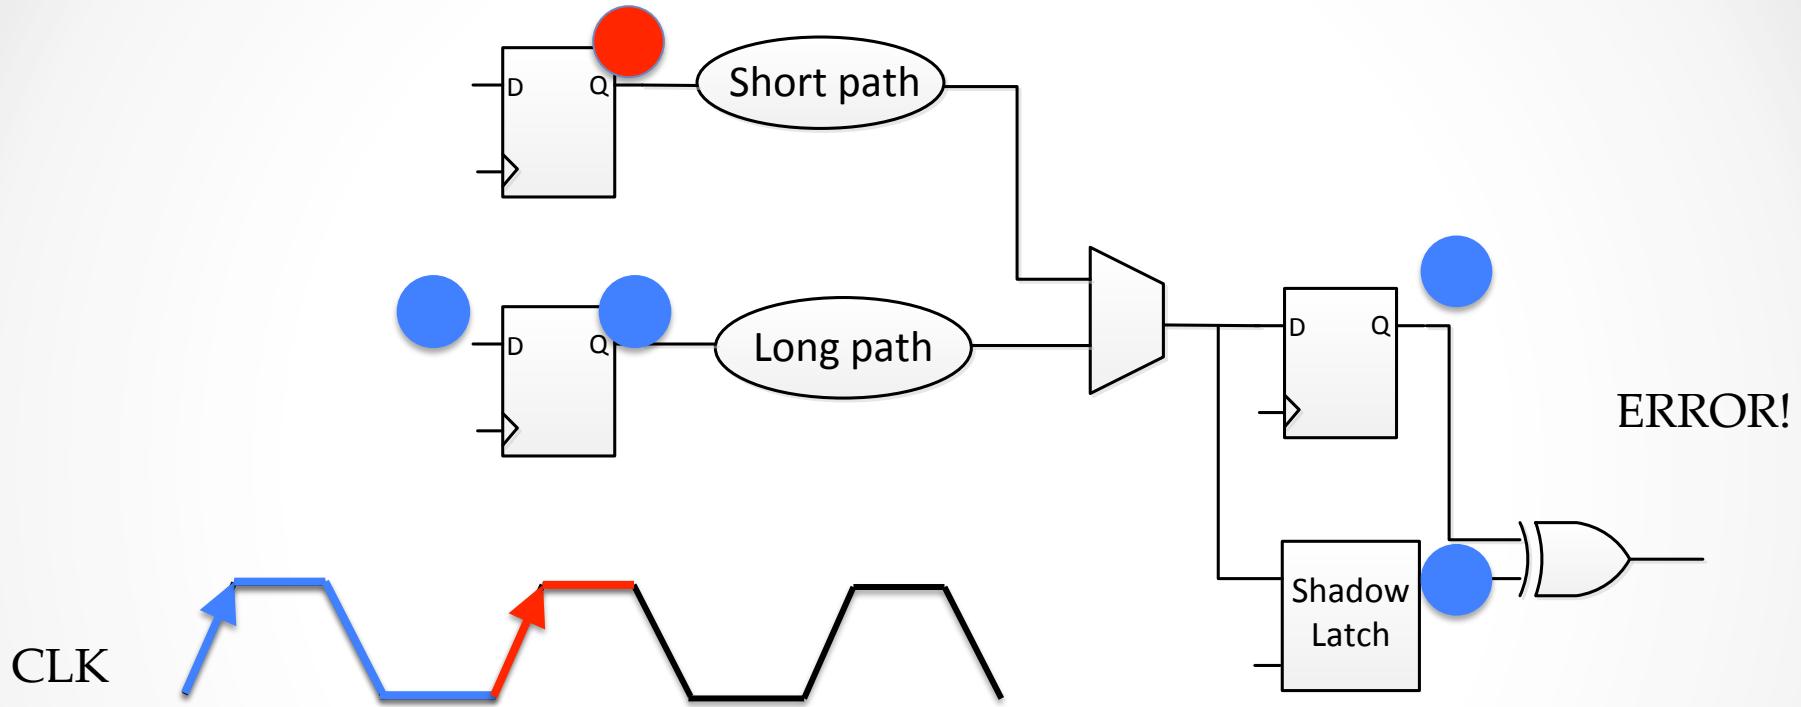
[3] D. Ernst, et al., *MICRO*, 2003.

Recovery mechanism



- Correct data restored in the following clock cycle
- All previous pipeline stages have to be flushed
- Program counter resumes at the next instruction

Razor – issues

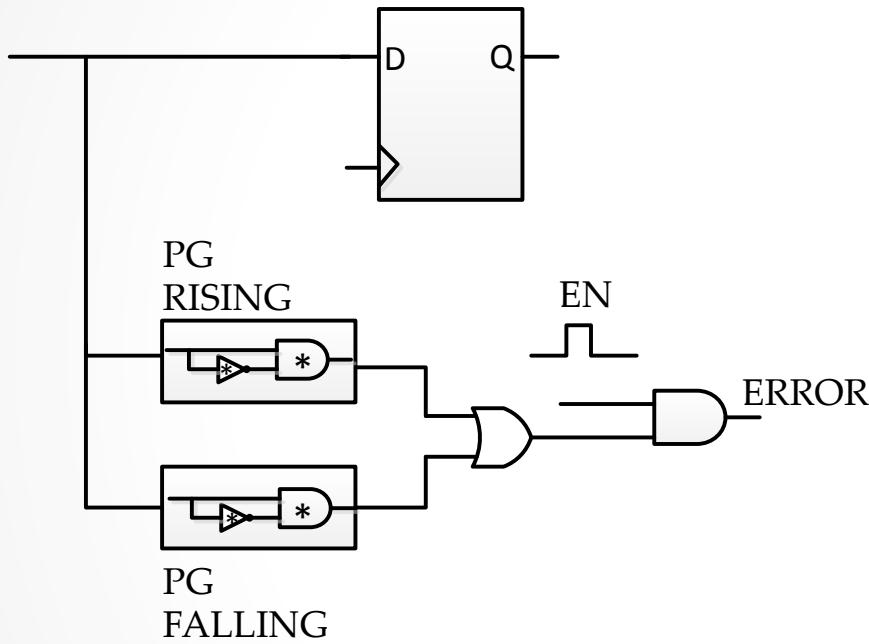


Good idea but has a lot of issues:

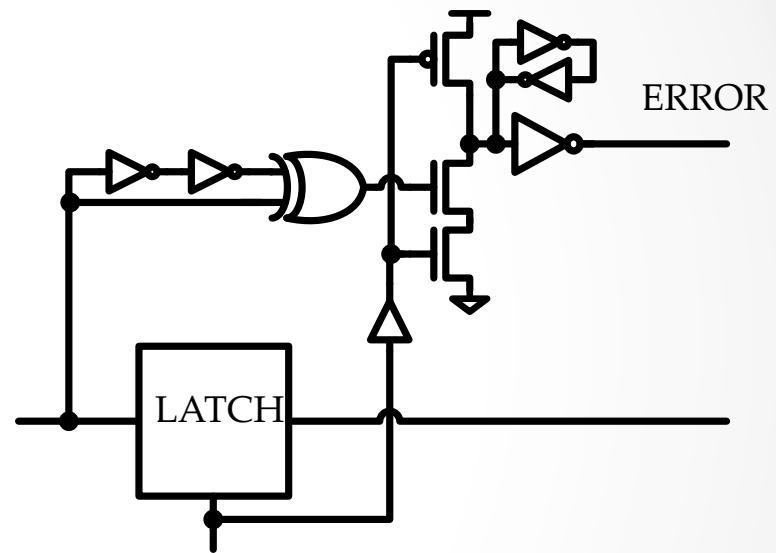
- Duty cycle constraints the shortest path
- Metastability at the output of the main FF
- The longest paths changed if a latch is added to FF

Improvements

ARM'10



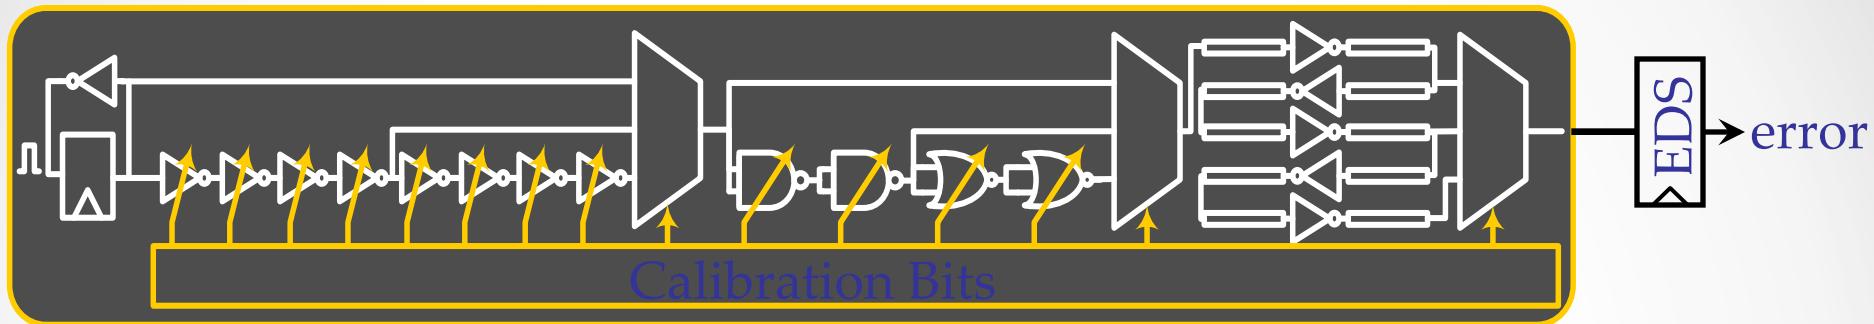
Intel'09



- ✓ No added buffers for short paths
- ✓ Data metastability detected w/o additional detector
- Detection window tuned in prefab

- ✓ Lower clock power
- ✓ Data metastability solved
 - Additional buffers needed for short paths

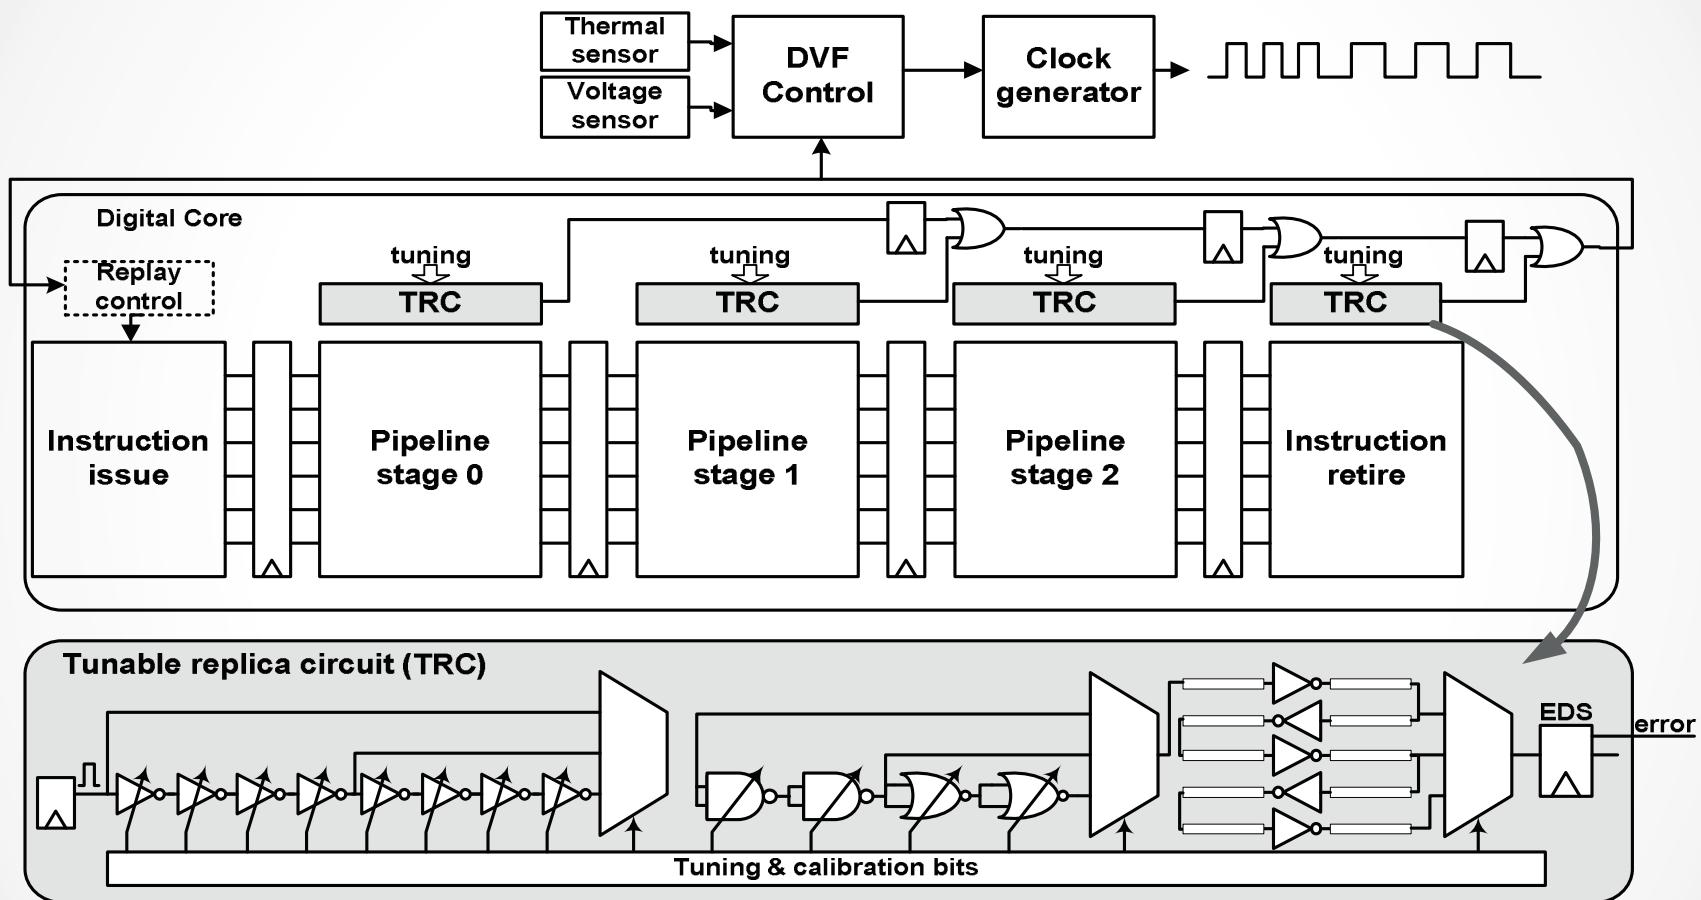
Tunable Replica Circuit (TRC)



Logic stages:

- Inverter
- NAND
- NOR
- Pass gates
- Repeated interconnects

TRC – cont'd



- Not so accurate as EDS, but no interfering with the longest path
- If properly tuned, no recovery mechanism needed

Observation circuits - summary

Advantages:

- ✓ Reduce margins for process variations
- ✓ EDS:
 - enable instruction recovery
 - detect errors in pipeline stages
- ✓ TRC:
 - capture clock-to-data delay per pipeline stage
 - have low design overhead

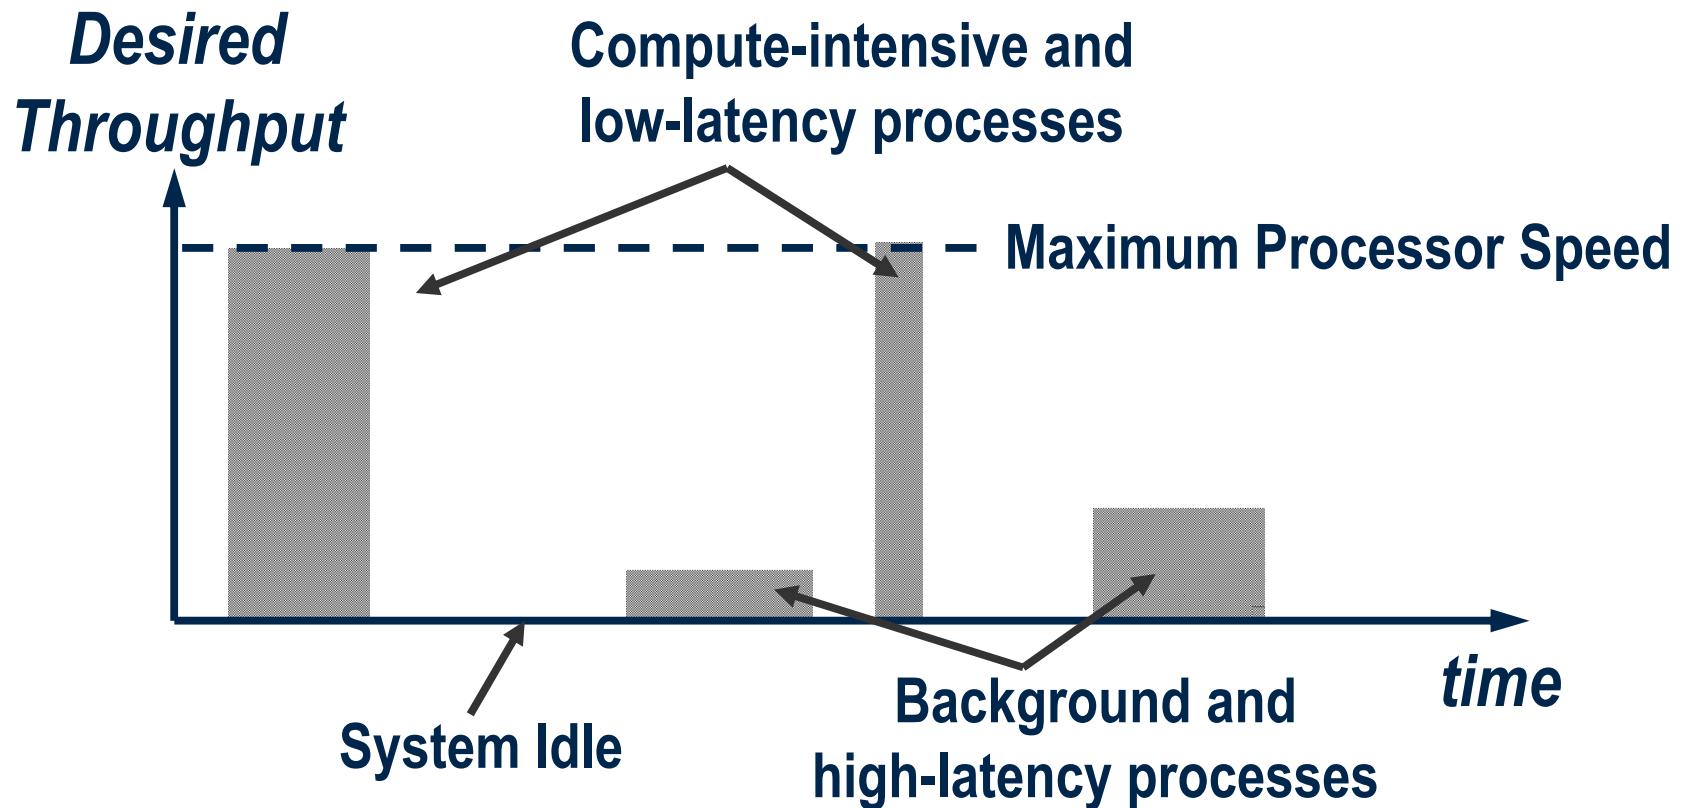
Disadvantages:

- EDS:
 - Adding buffers for short paths
 - Metastability issues
 - The longest path is affected by additional circuits
- TRC:
 - Cannot detect local dynamic variations
 - Requires margin between TRC & the longest path
 - Requires post-silicon calibration

Overview

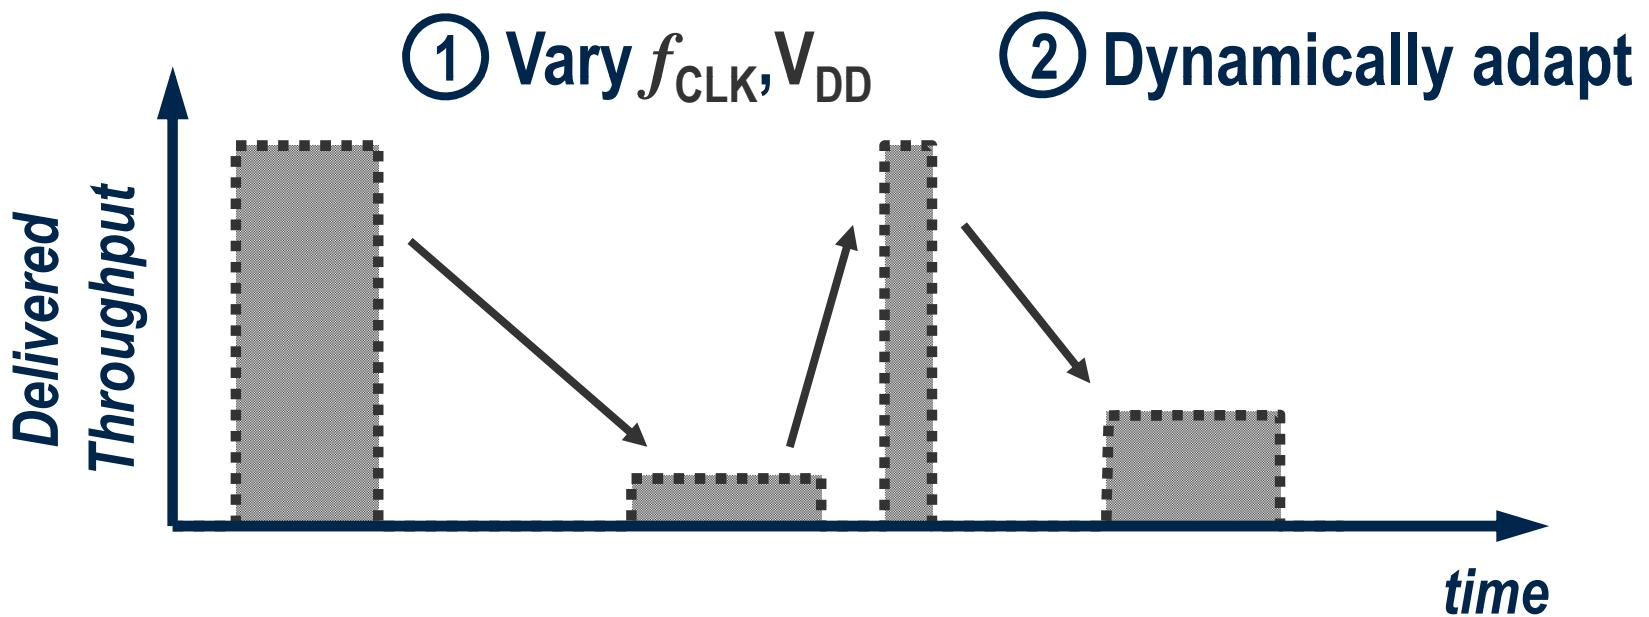
- Error detection and correction circuits
- Dynamic voltage and frequency scaling
- Control voltage through DC-DC converters
- Conclusions

DVFS



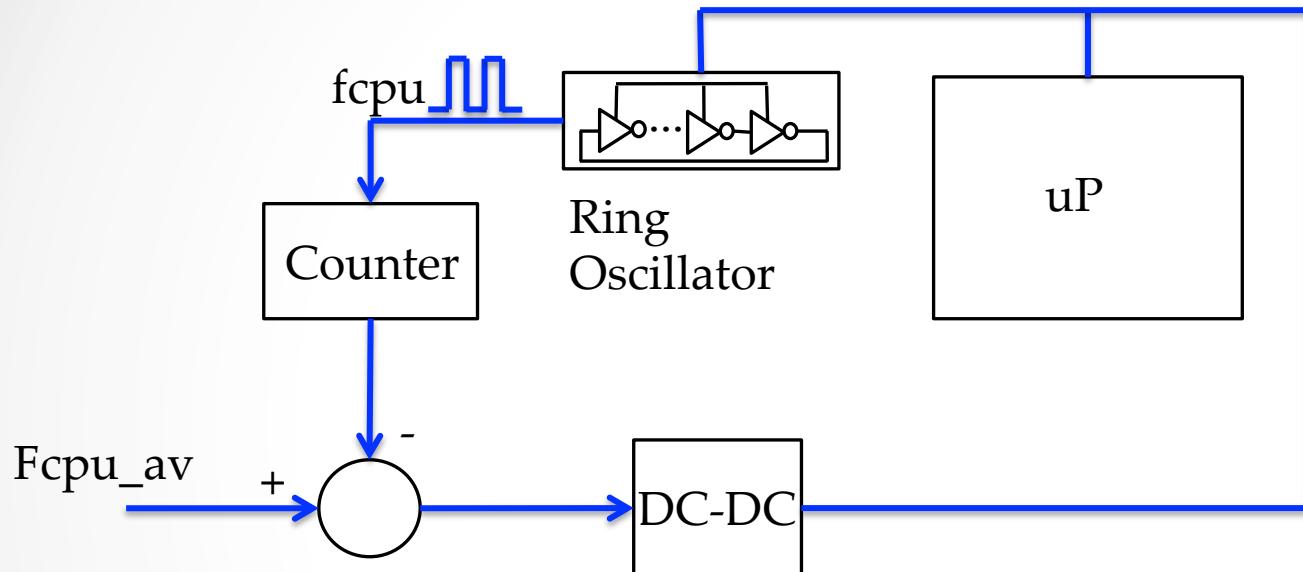
Burd, ISSCC'00

DVFS – cont'd



Burd, ISSCC'00

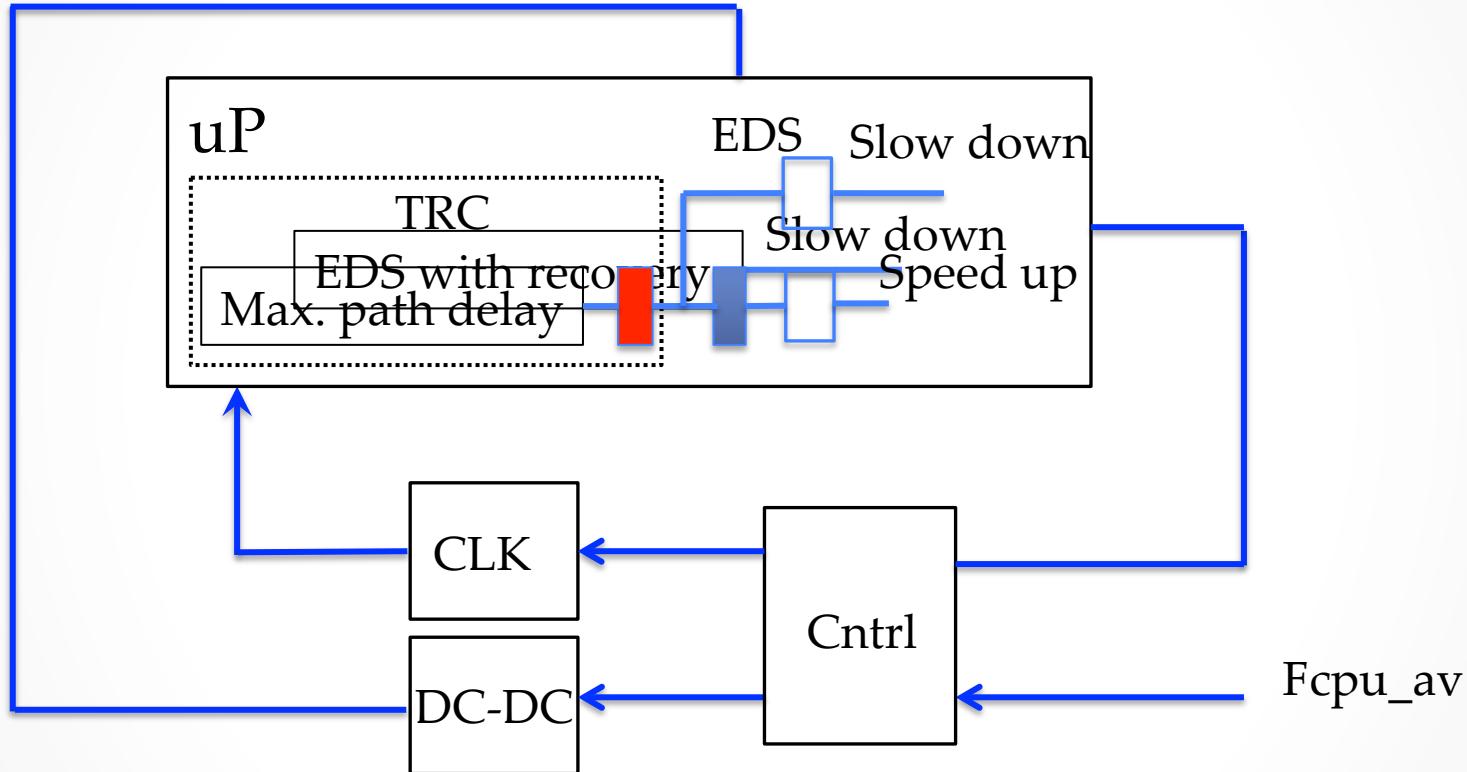
Traditional DVFS



Burd, ISSCC'00

- Traditional DVFS use a ring oscillator for frequency detection (Xscale, PowerPC, Pentium M, ...)
- Impossible to use nowadays: ring oscillator frequency change does not reflect the cpu frequency change
- Different paths have different behavior with Vdd change

Resilient DVFS



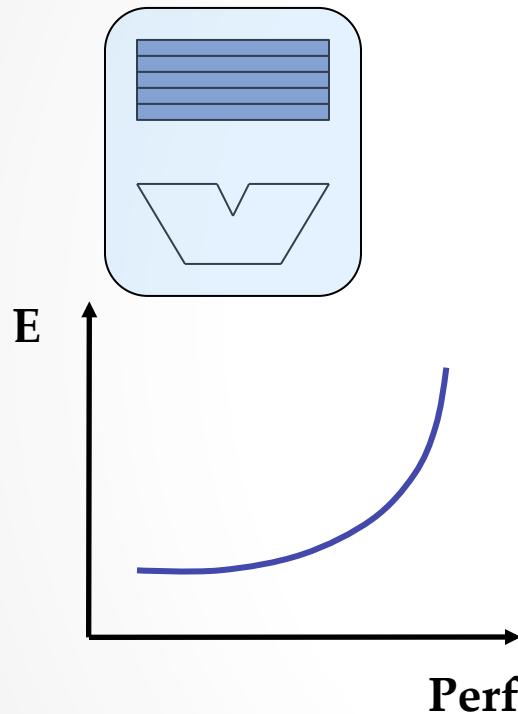
- Resilient DVFS use TRCs and EDS as freq. indicator
- Two options: TRC+EDS or EDS with recovery

RAVEN microprocessor

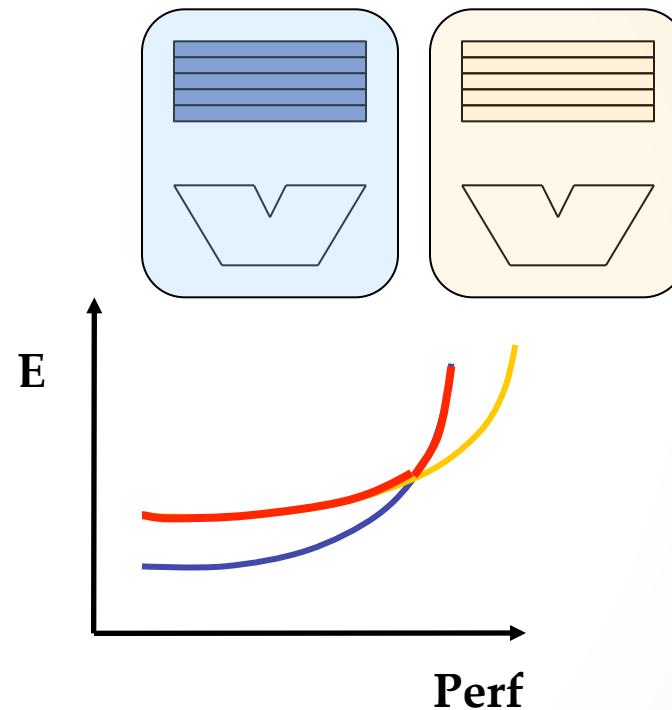
- Implemented in the newest technology: 28nm!
(apple i5/i7 cores in 32nm, altera chips in 28nm)
- Mobile applications: battery life important!
- Manycore architecture: exploiting parallelism!
- Error detection circuits for observation:
improvement over ARM architecture!
- Unconventional DVFS scheme

Motivation - DVFS

- Single core



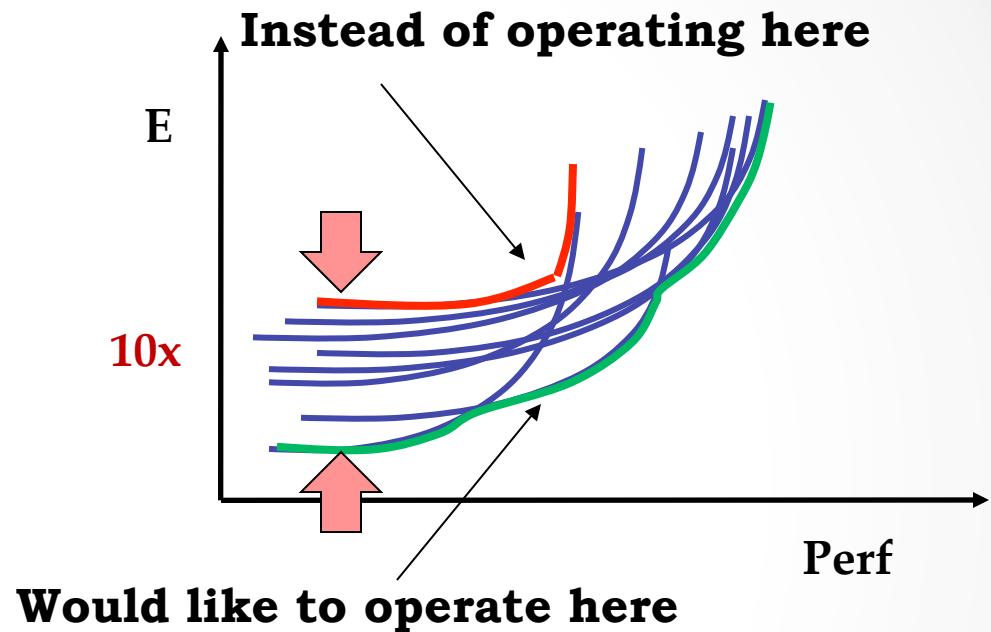
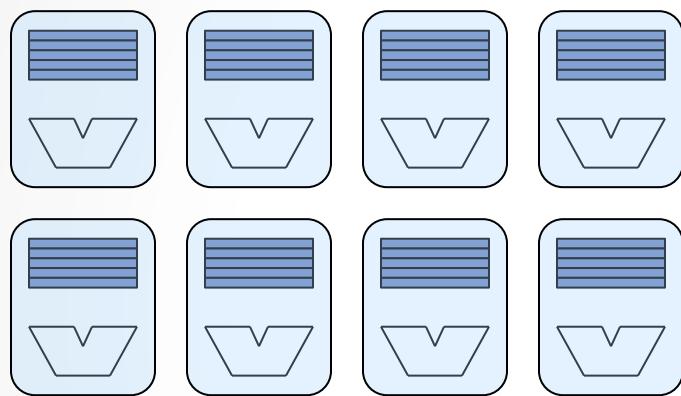
- Two cores



- Performance dictated by the slower core
 - In a conventional DVFS synchronous system

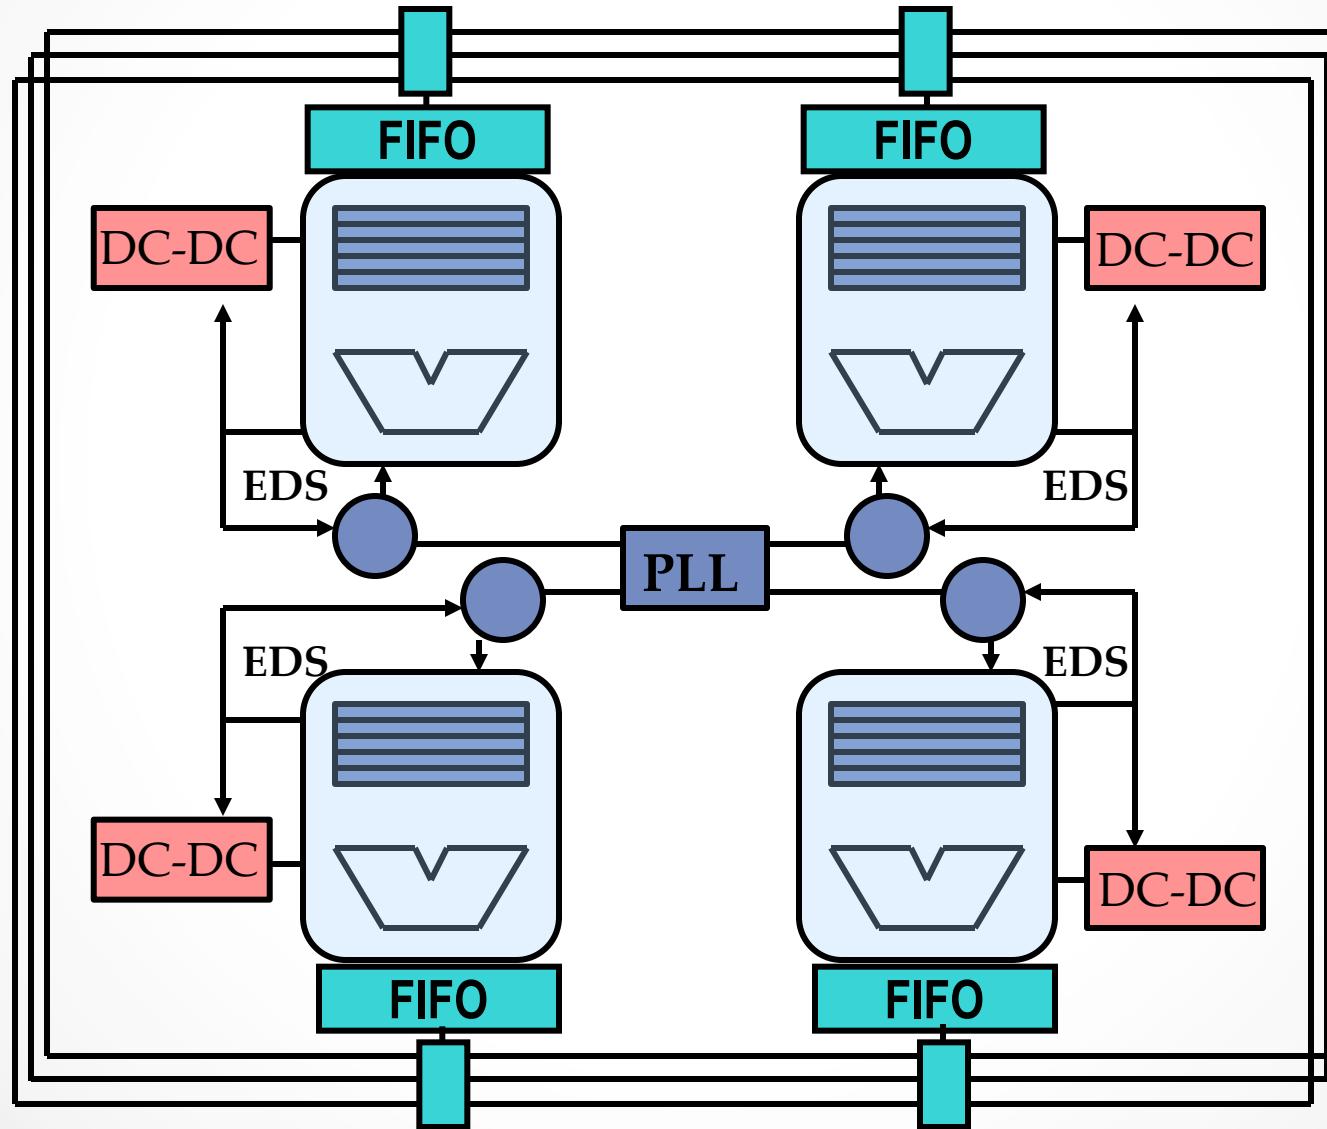
Goal

- Manycore

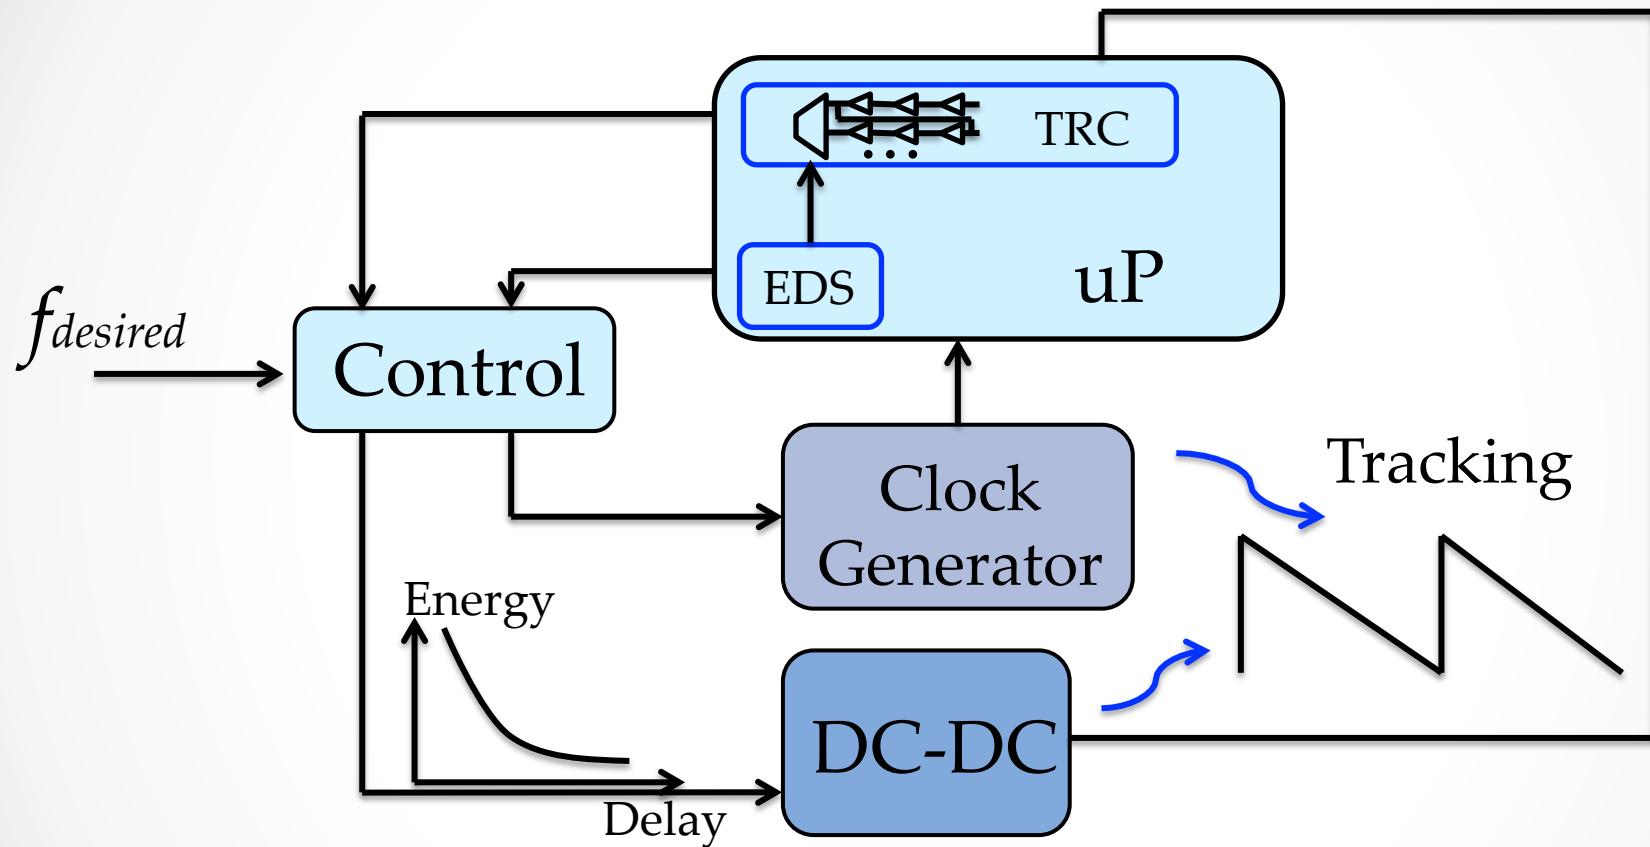


- Make sure to operate at the optimal energy-performance point

Per-Core Supply/Clock Control



DVFS on manycore processor

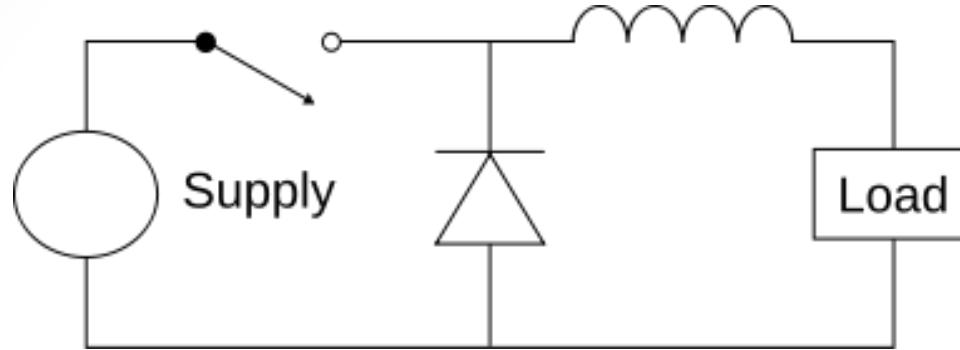


- Allow the ripple at the output and track the voltage through clock generation for better energy-efficiency

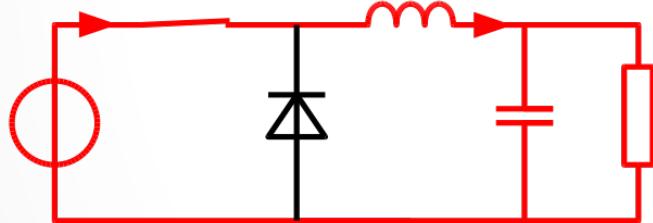
Overview

- Error detection and correction circuits
- Dynamic voltage and frequency scaling
- Control voltage through DC-DC converters
- Conclusions

DC-DC converters



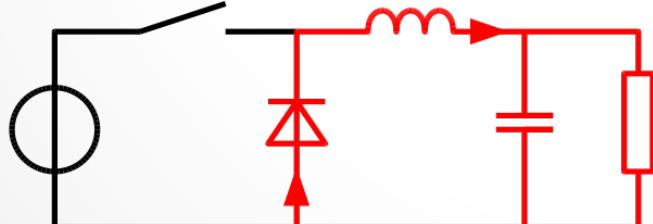
-State



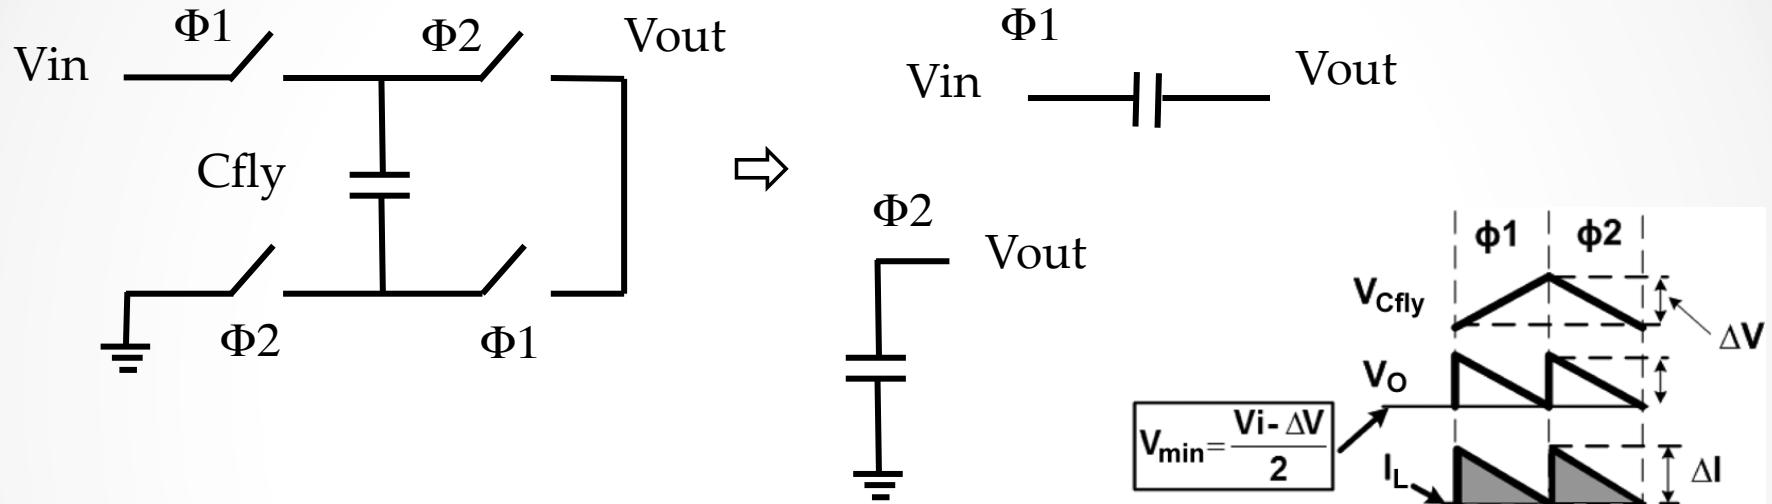
Two phases:

1. Loading the energy from the battery
2. Transferring loaded energy to the output

-State

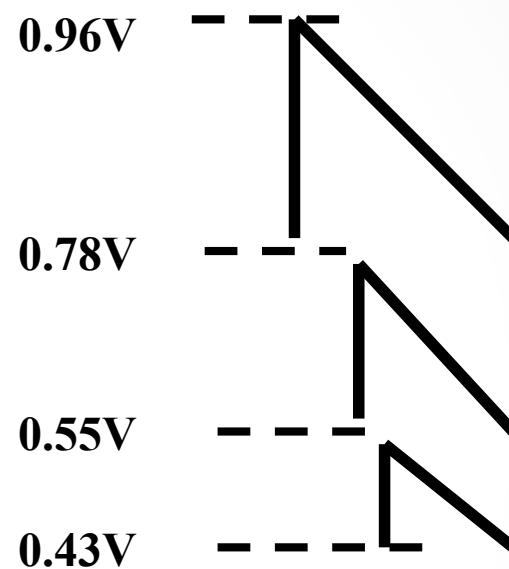
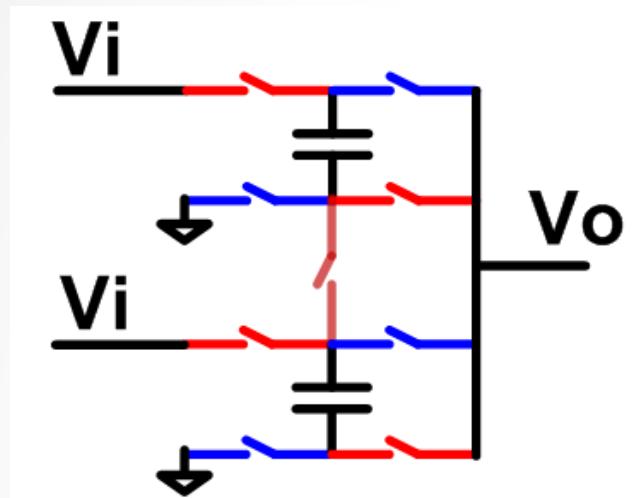


Switched Capacitor DC-DC Converter



- Advantages:
 - Fully integrated on a chip
 - Smaller area (no inductive components)
 - Large power density
- Disadvantages:
 - Discrete output voltage
 - Low efficiency

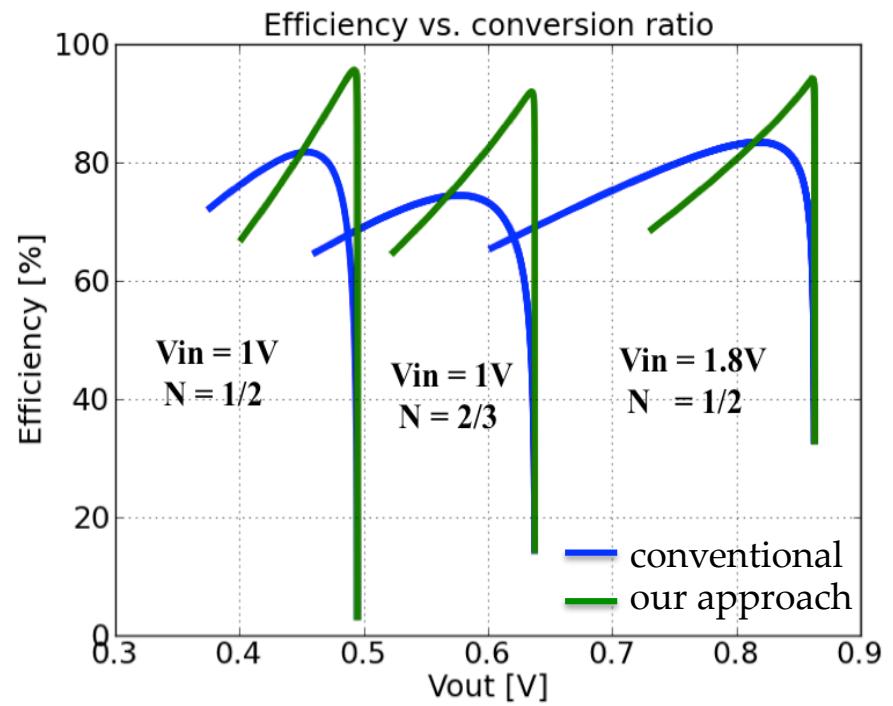
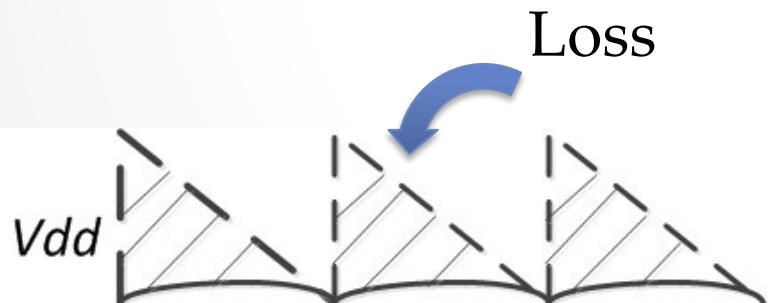
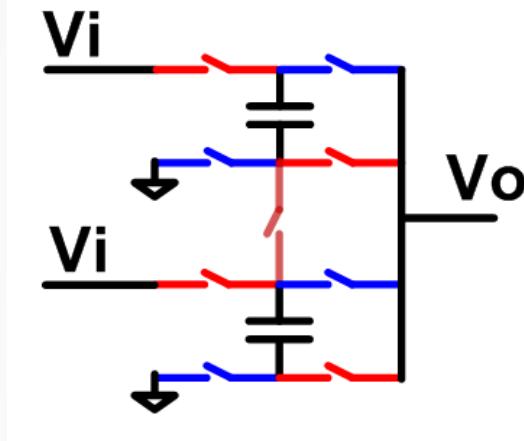
Discrete points solution



H-P. Le, ISSCC 2011

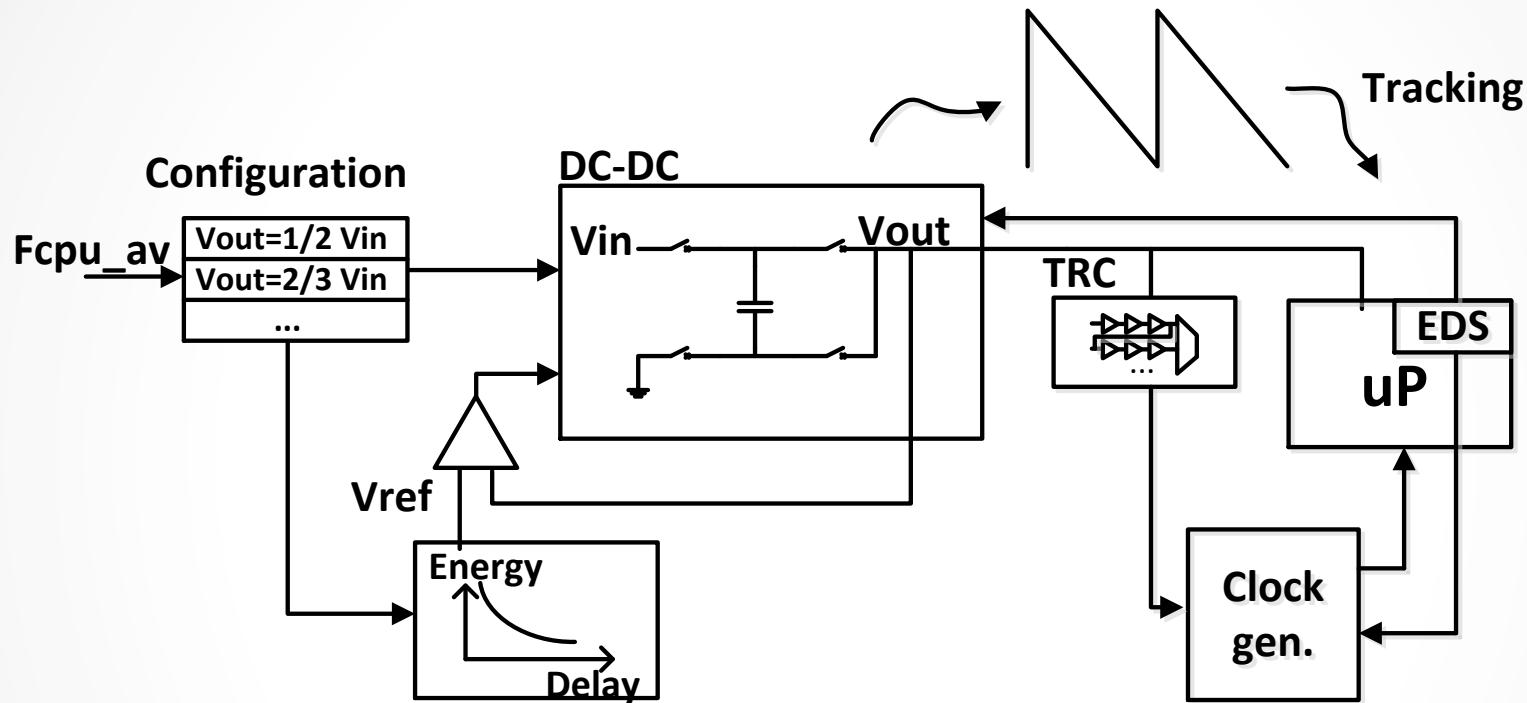
V_{in} [V]	ratio	V_{out} [V]	Range[V]
1	1/2	0.5	0.45 – 0.55
1	2/3	0.67	0.55 – 0.71
1	1/3	Too small	Too small
1.8	1/2	0.9	0.79 – 0.96

Efficiency solution

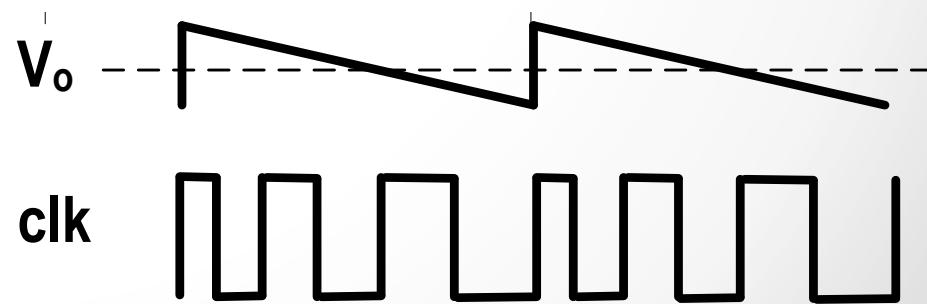


Efficiency can be improved
by more than 10%!

DVFS details



- Set the ripple size at the optimal energy-delay point



Overview

- Error detection and correction circuits
- Dynamic voltage and frequency scaling
- Control voltage through DC-DC converters
- Conclusions

Summary

- Process variations introduce difficulties in circuit design
- Power is a critical design constraint
- Observation circuits needed in order to avoid too conservative design decisions (EDS and TRC)
- Multicore/manycore architectures open spatial dimension for energy optimization through DVFS
- Fine granularity V-F control enabled through performance observation and DC-DC converters

Acknowledgement



Marie Curie FP7 People program

