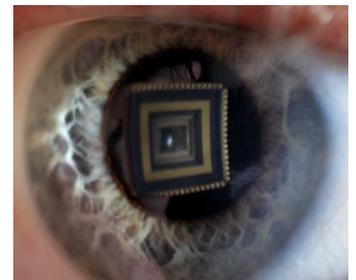


Energy-Modulated Computing

Alex Yakovlev
School of EECE
Newcastle University

EACO 19 October 2011



Outline

- **Introduction: better resource awareness**
- **Energy-uncertainty-QoS interplay**
- **Power-adaptive computing**
 - **System design aspects**
 - **Power proportional and power efficient systems**
 - **Experiment in power-proportionality**
 - **System Design for energy-harvester power supply**
 - **Circuits for power-adaptive systems**
 - **SRAM**
 - **Voltage sensors**
 - **Novel power electronics**
- **Conclusion**

Introduction

- **Systems that are performant and resourceful**



- **Green, energy-frugal, power-proportional**
- **Energy supply perspective: battery, harvester, power control, regulation**
- **Energy consumption spectrum: Mwatts (data plants), 100s Watts (many core chips), uWatts (implantable devices)**

Introduction

Challenge: does energy aspect modulate computations?

- both in every system design case,
- as well as an evolution process (e.g. can we come up with some “Mooreish” law of system performance driven by energy levels?)

Before the mountain assault:

- Measure your terrain – identify important constraints and criteria
- Prepare your gear – work from design examples



Introduction

Working conditions:

Energy-constrained systems

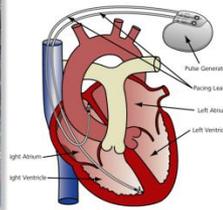
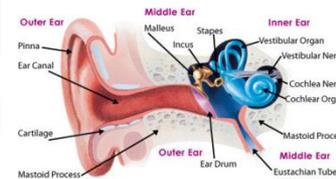
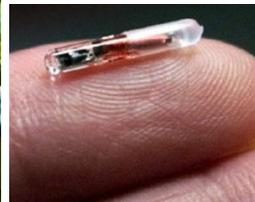
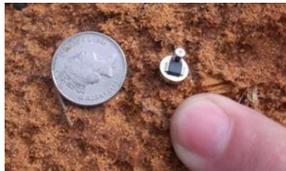
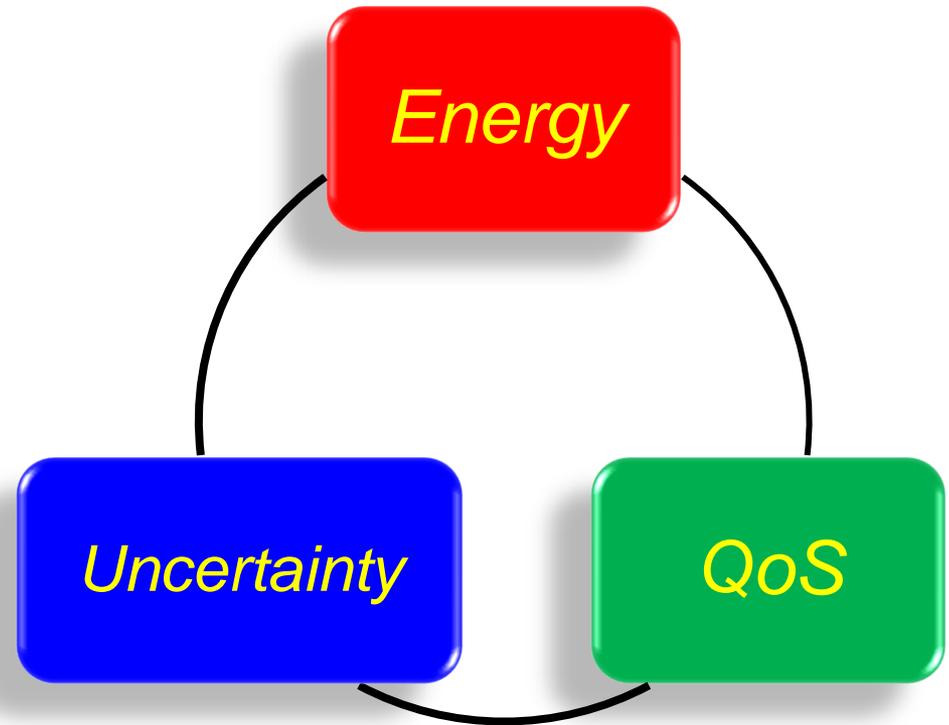
- Solar energy, e-beam power supply, small batteries, ...

Unreliable power supply

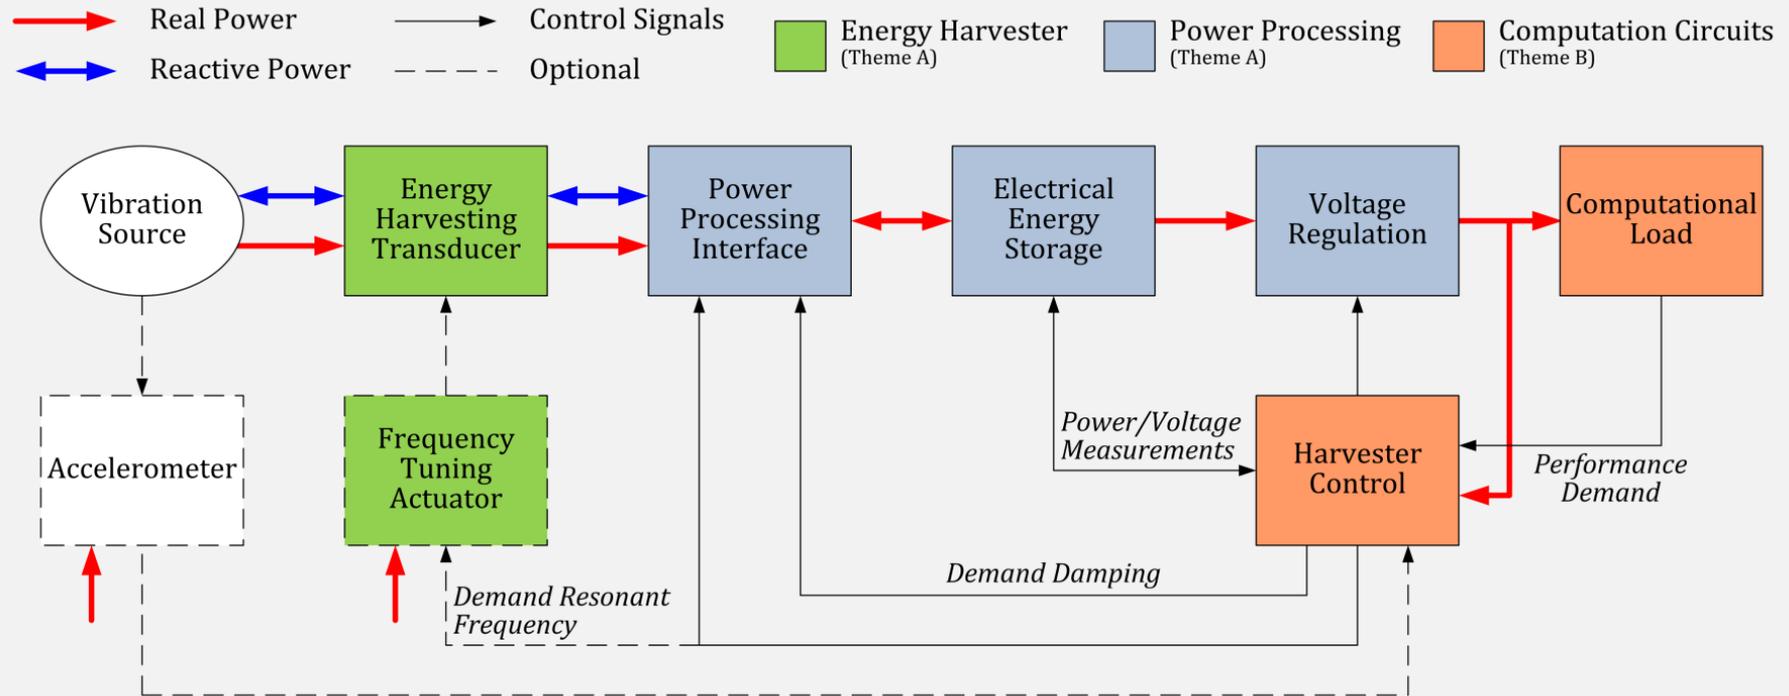
- Voltage fluctuations, low battery, ...

Hostile environments

- High/low temperatures, noise, ...

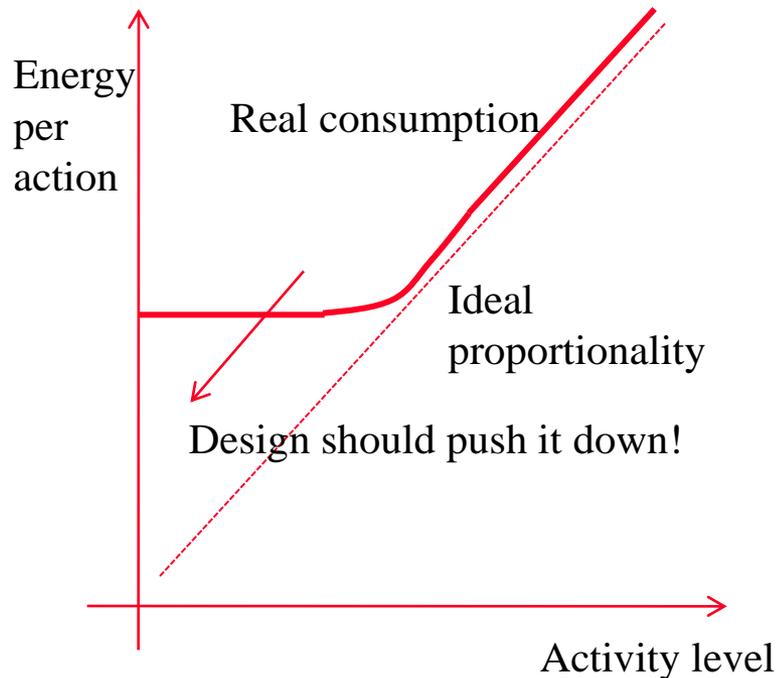


The “Holistic” Project



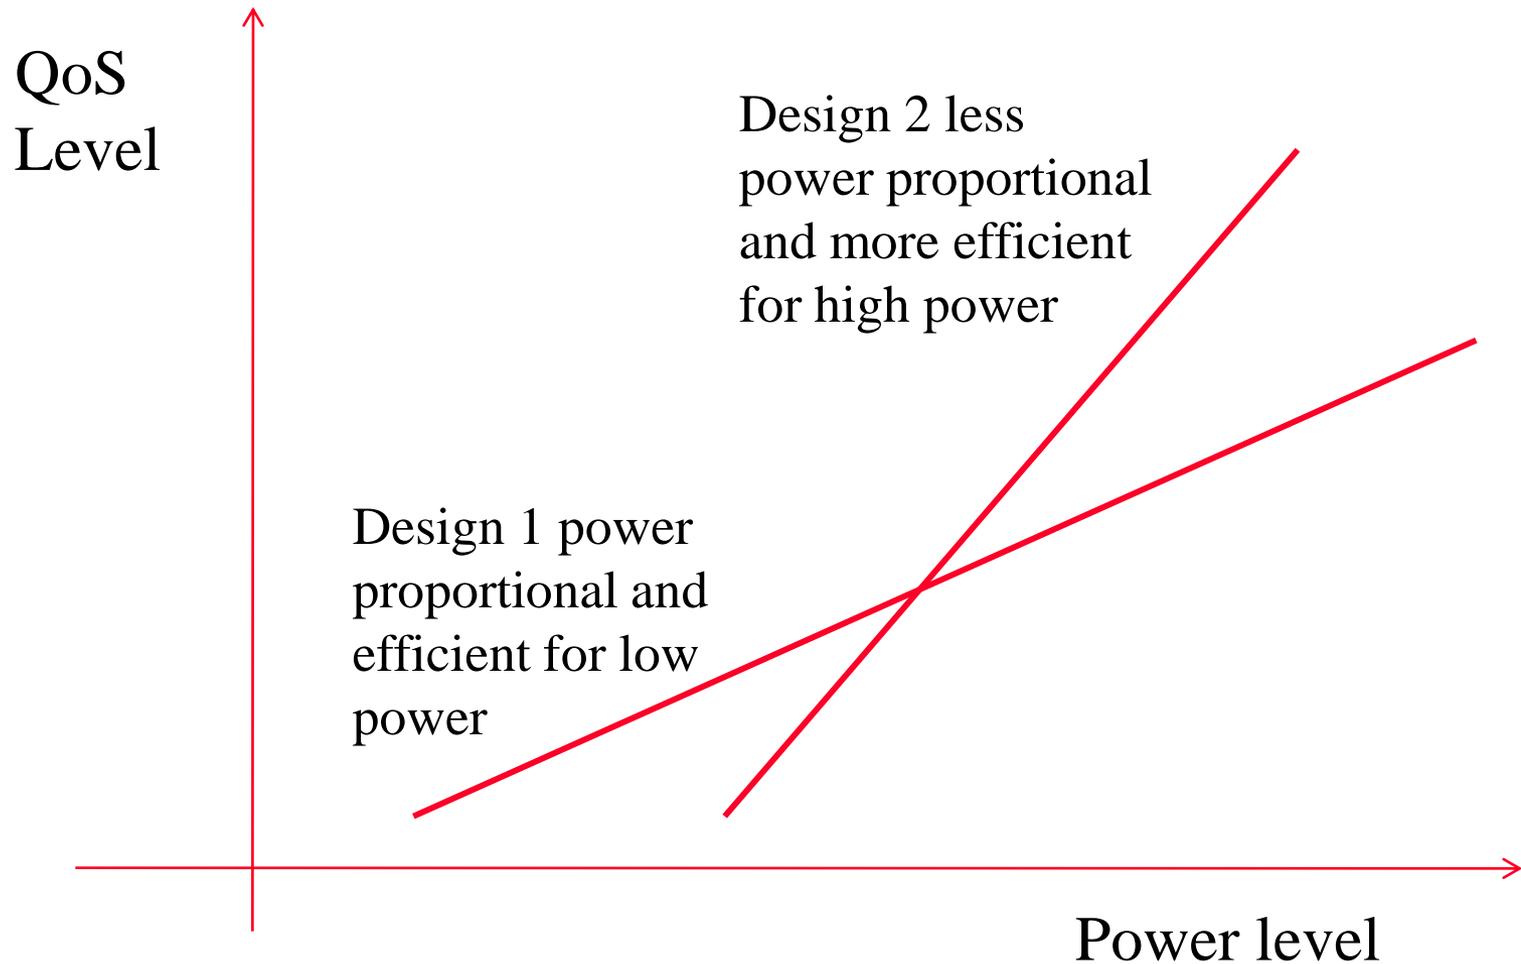
EPSRC project: Next Generation Energy-Harvesting Electronics: Holistic Approach, Universities of Southampton, Bristol, Newcastle and Imperial College London
 "Energy Harvesting Systems: A Block Diagram (2010, July 16). Holistic Energy Harvesting [Online]. Available: <http://www.holistic.ecs.soton.ac.uk/res/eh-system.php>"

Energy-proportional computing

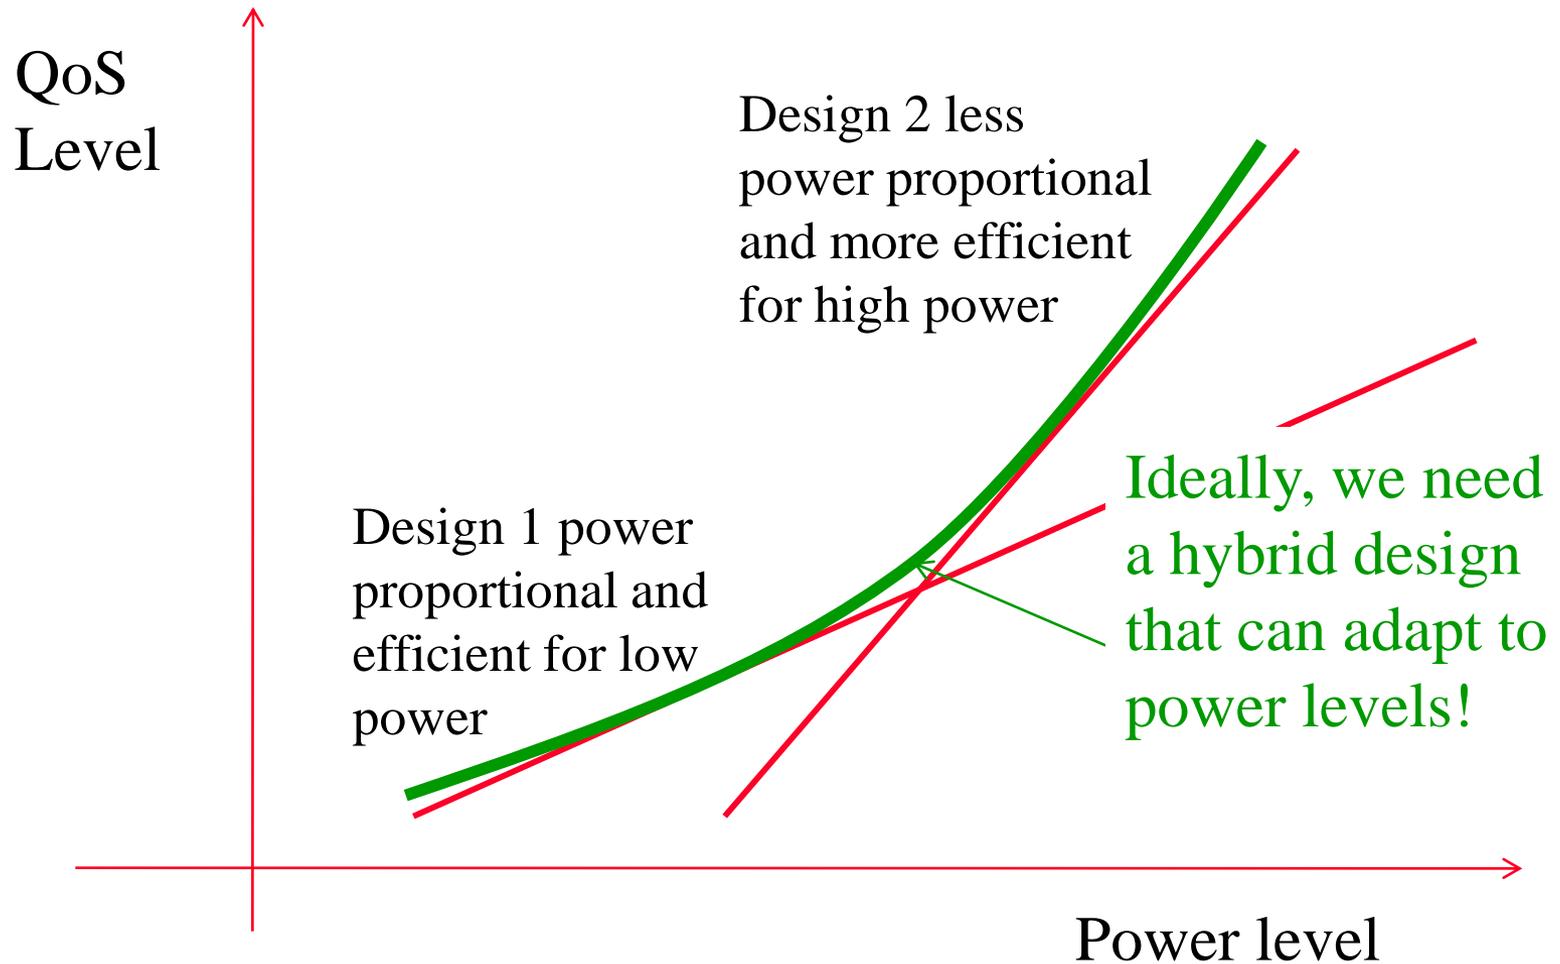


“Systems tend to be designed and optimized for peak performance. In reality, most computation nodes, networks and storage devices typically operate at a fraction of the maximum load, and do this with surprisingly low energy efficiency. If we could **design systems that do nothing well** (as phrased by David Culler), major energy savings would be enabled. Accomplishing **energy-proportional computing** requires a full-fledged top-down and bottom-up approach to the design of IT systems.” (from Jan Rabaey’s lecture **The Art of Green Design: Doing Nothing Well** – March 2010)

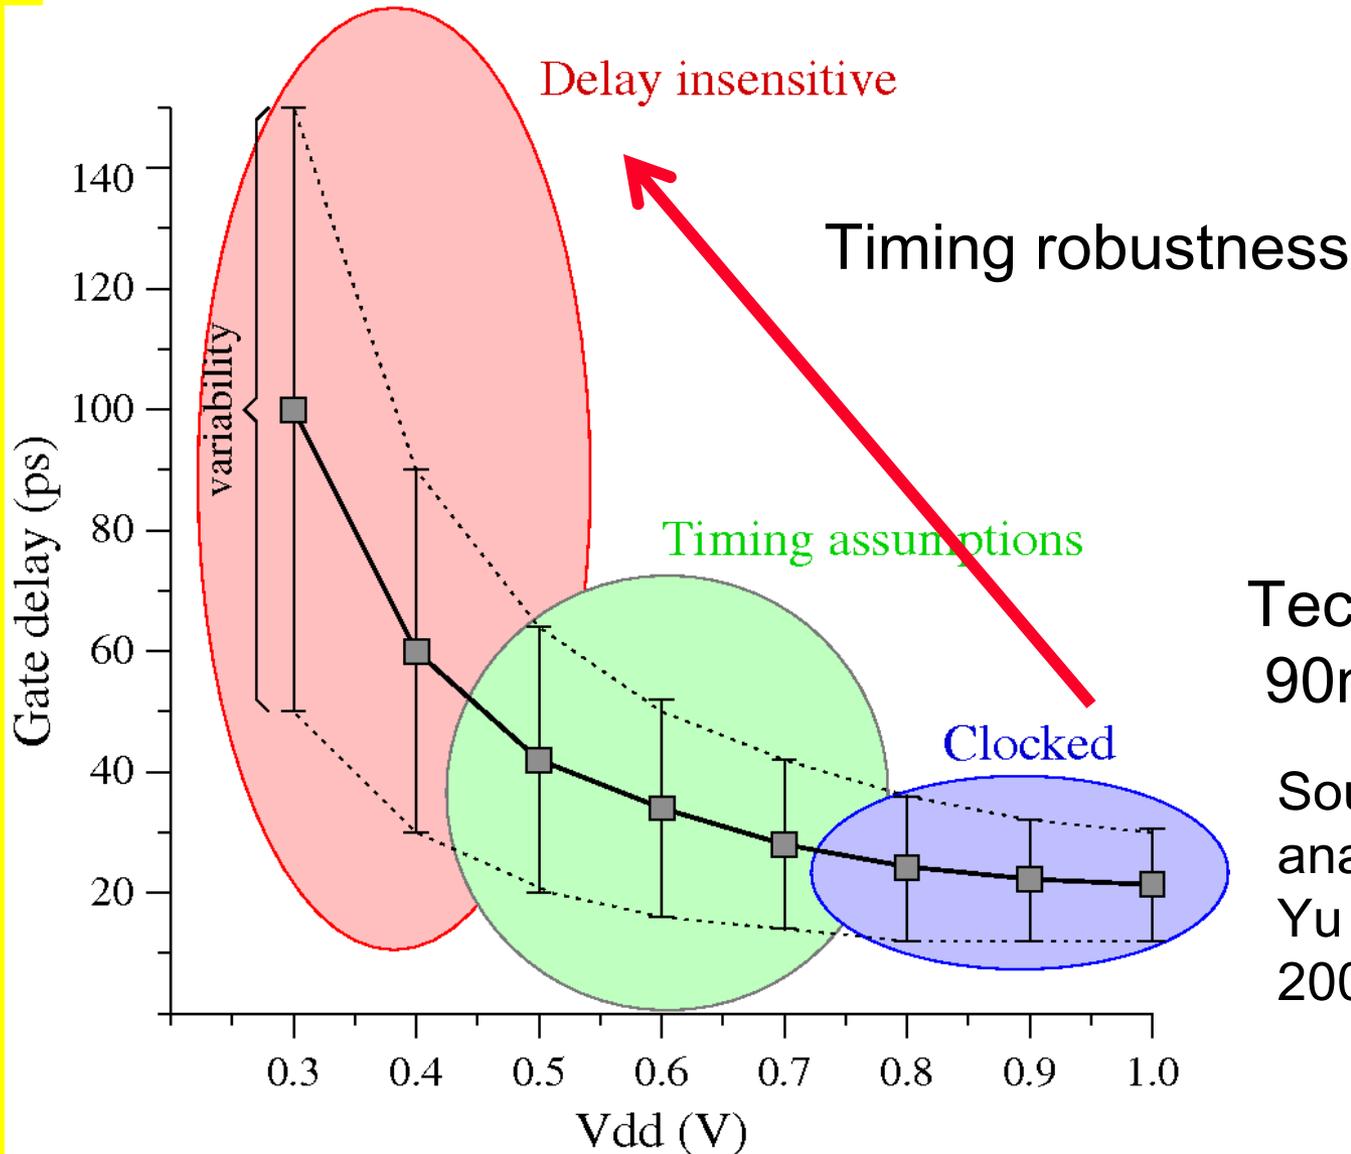
Power-proportional vs Power-Efficient



Power-proportional vs Power-Efficient



Relationship with timing variability



Technology node:
90nm

Source of variability
analysis:
Yu Cao, Clark, L.T.,
2007

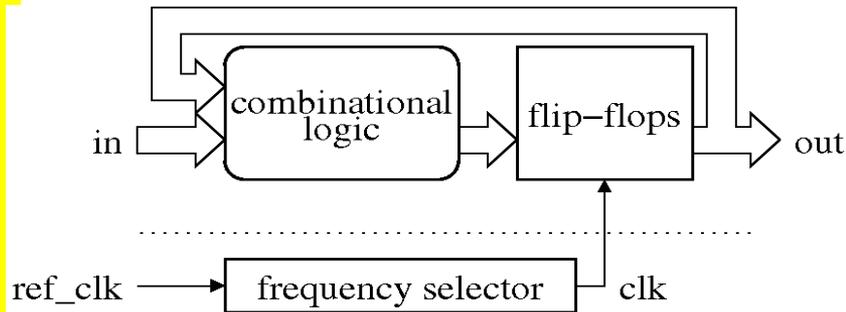
Power proportionality experiments

- **Experiment setup**
 - **Synthesised for Faraday library**
 - **Based on UMC 90nm technology process**
- **8-bit Booth's multiplier – speed as the QoS**
 - **SPECTRE analogue simulation**
 - **Runs at 1V, 0.9V, ..., 0.1V source voltage**
- **Iterative FFT – precision as the QoS**
 - **VCS and PrimeTime-PX digital simulation**
 - **Runs as nominal 1V source voltage**

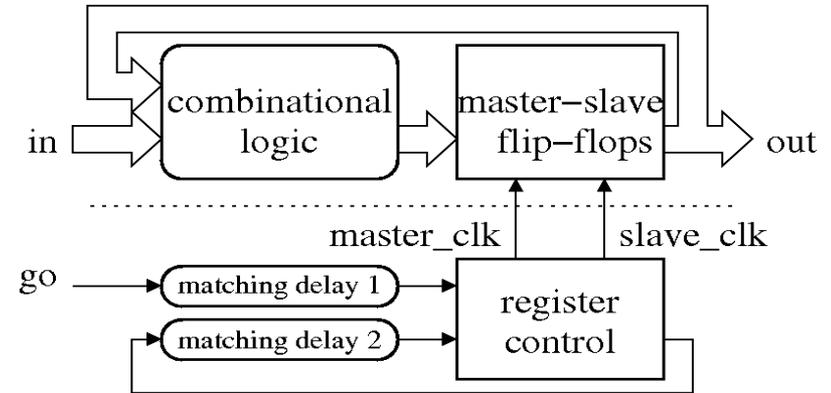
Benchmarks: 8-bit Booth's Multiplier

- **Synchronous**
 - **Rigid 1GHz clock**
- **Frequency scaling**
 - **Tuned for 1GHz, 500MHz and 250MHz**
- **Asynchronous, bundled data**
 - **Extra control logic and delay lines**
- **Asynchronous, dual-rail**
 - **Double comb. logic and FF size (more leakage)**
 - **Extra completion detection and single-rail to dual-rail converters**
 - **Double switching activity (spacer/code-word)**

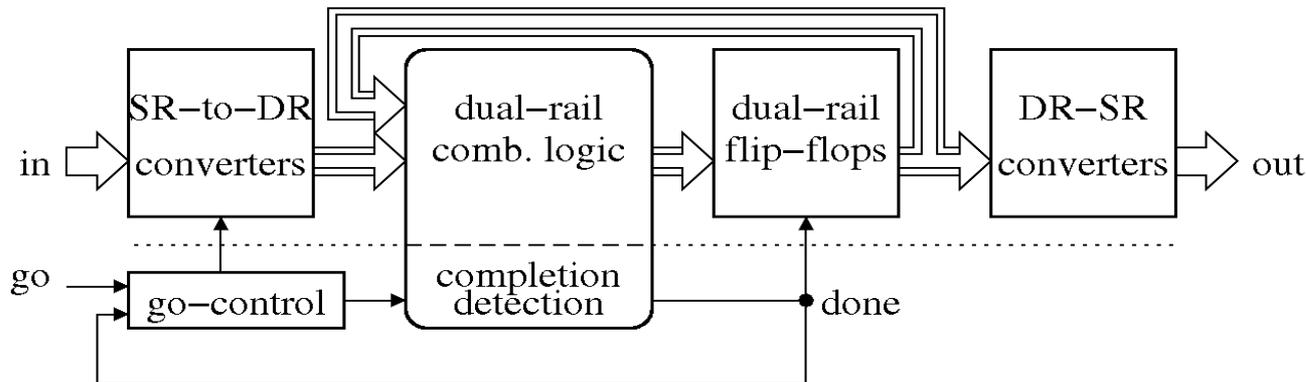
Benchmark Architectures



Adaptive frequency scaling



Bundled data



Dual-rail

Multiplier: Simulation Results

	Area (cells)	Energy consumption (pJ) @ Computation time (ns)								
		1.0V	0.9V	0.8V	0.7V	0.6V	0.5V	0.4V	0.3V	0.2V
A	2495	22.5 @18.0	17.9 @18.0	13.8 @18.0	-	-	-	-	-	-
B	2495	22.5 @18.0	17.9 @18.0	13.8 @18.0	10.4 @36.0	7.5 @36.0	5.2 @72.0	-	-	-
C	2931	39.6 @17.2	30.8 @19.2	23.3 @22.2	17.2 @26.9	12.3 @35.2	8.4 @52.1	5.4 @96.1	3.3 @263.9	-
D	7683	279.0 @38.1	217.3 @41.6	166.1 @49.5	123.8 @61.5	87.1 @106.1	60.8 @142.6	39.1 @283.6	24.3 @831.4	19.1 @4140

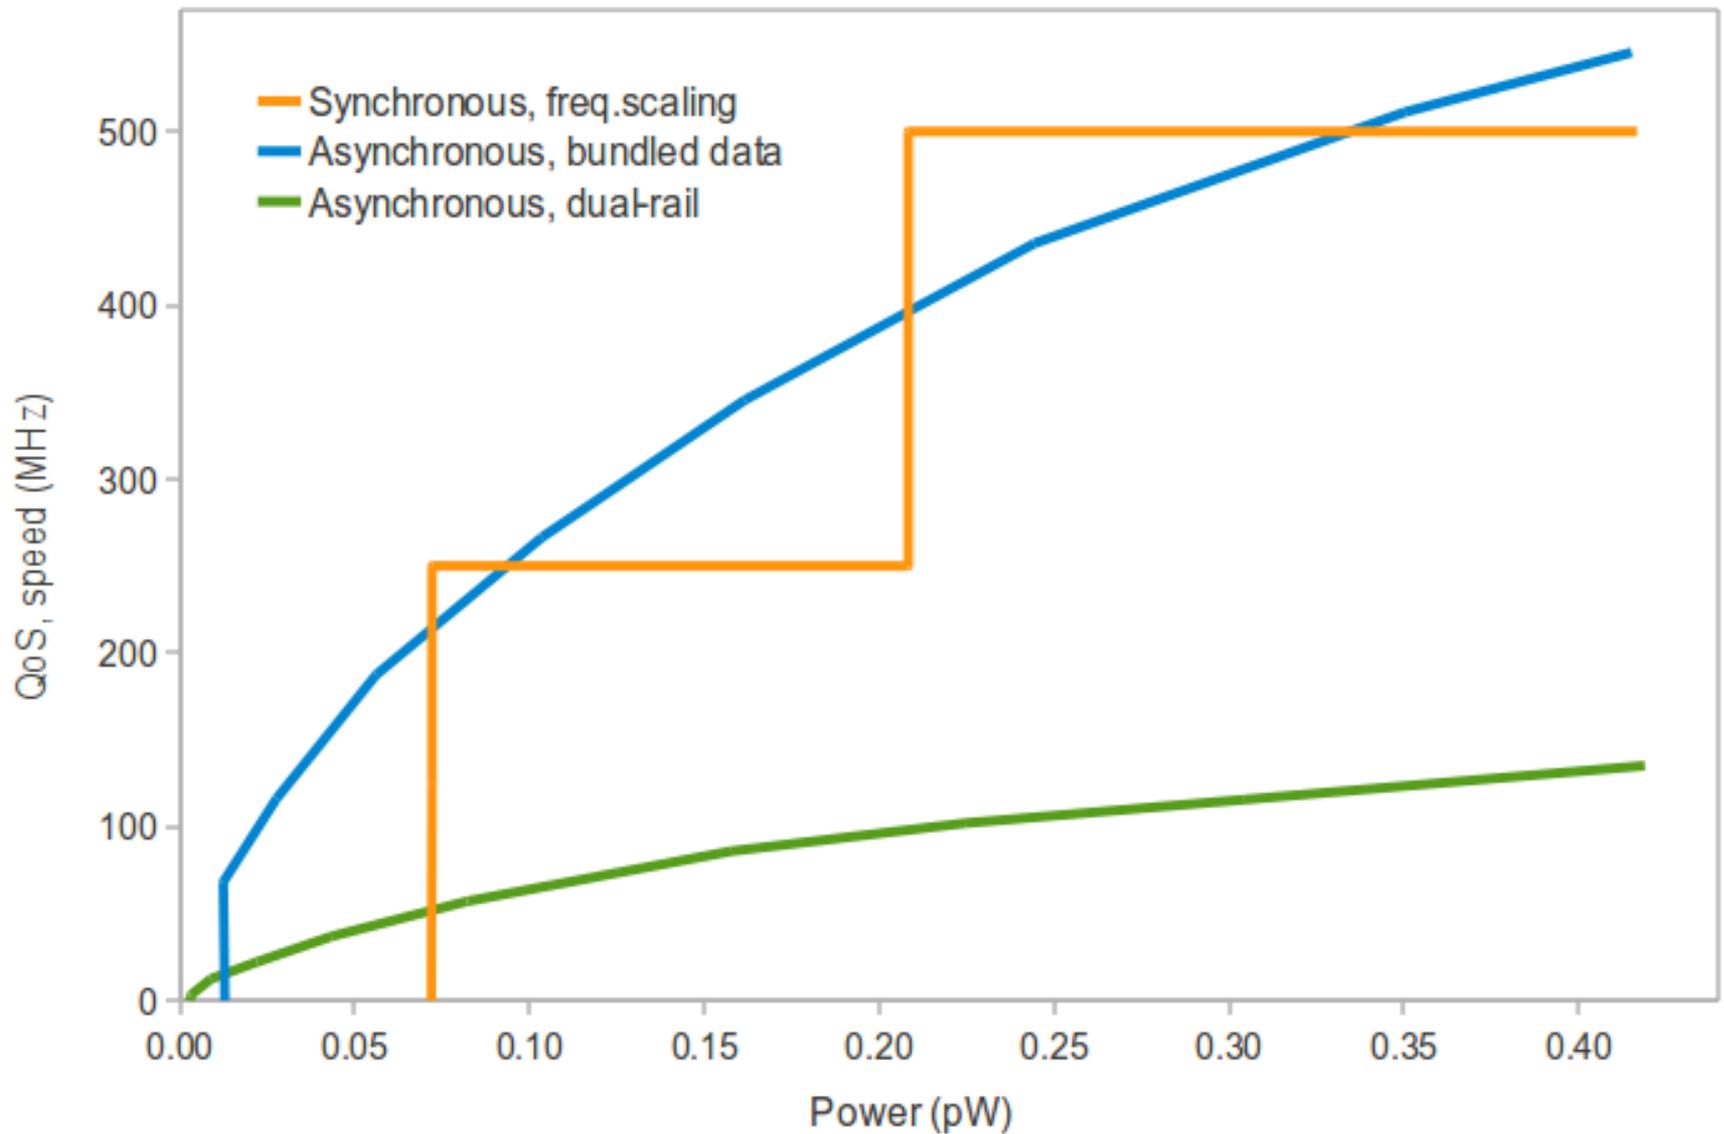
A - synchronous, fixed frequency @1GHz

B - synchronous, frequency scaling @1GHz, 500MHz, 250MHz

C - asynchronous, bundled data

D - asynchronous, dual-rail

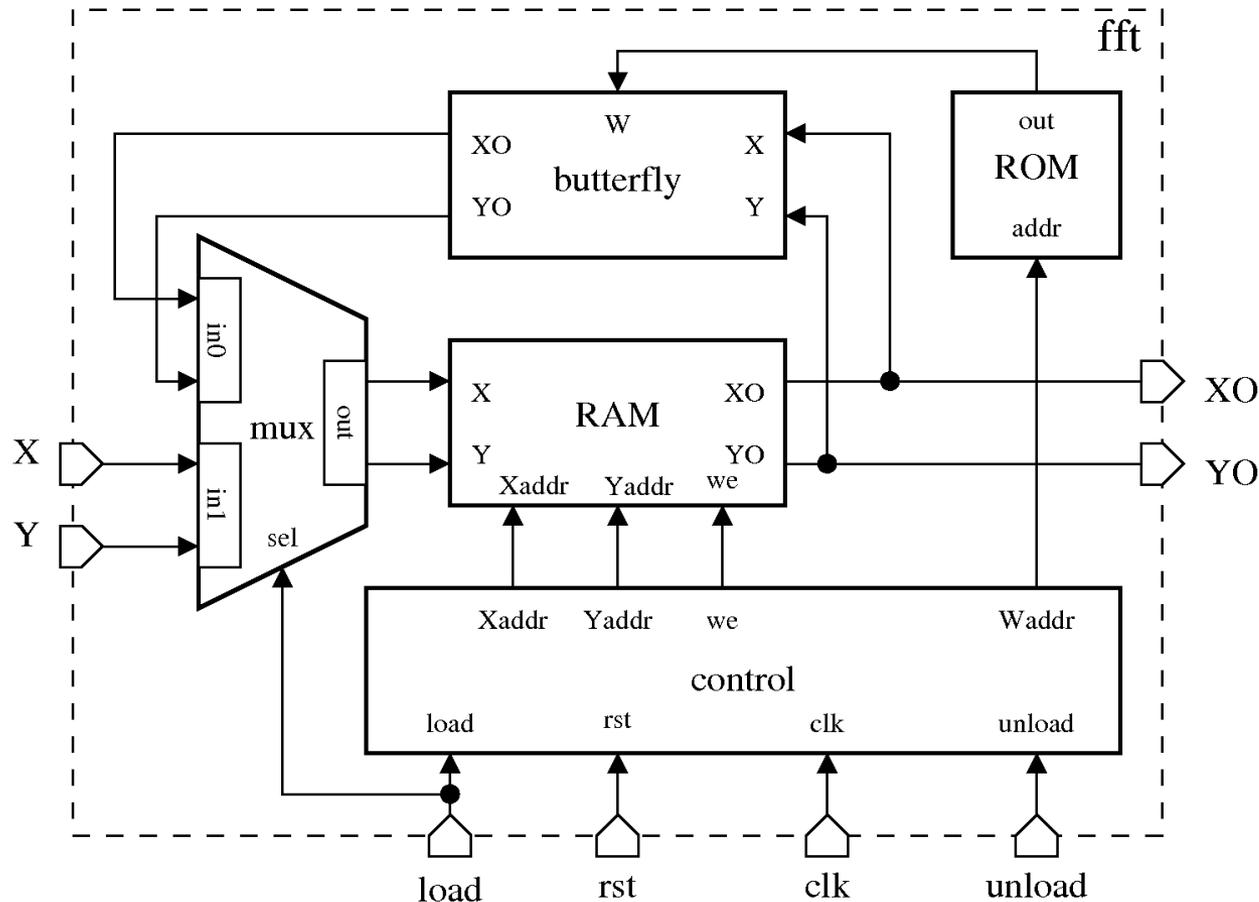
Multiplier: Quality of Service



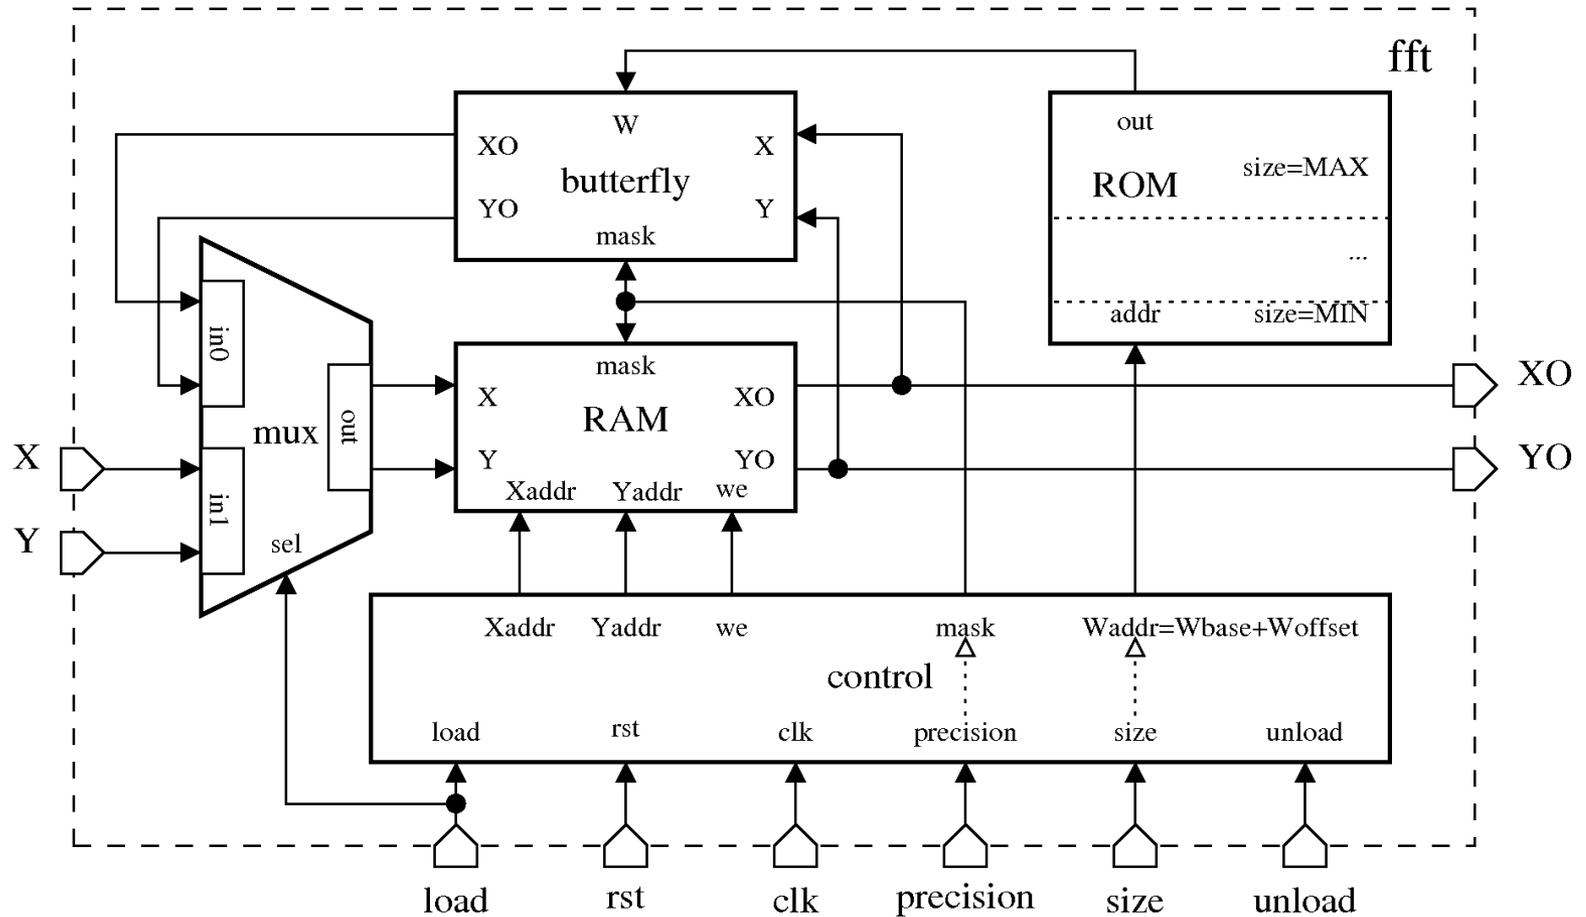
Benchmarks: Iterative FFT

- **Non-reconfigurable FFT**
 - **Fixed transform size (1024 points)**
- **Reconfigurable sample size**
 - **Changeable transform size (1024/512/256 points)**
- **Reconfigurable data precision**
 - **Variable data representation (16bit / 12 bit / 8 bit)**
- **Reconfigurable sample size and data precision**
- **Asynchronous reconfigurable FFT implementation**
 - **Bundled data with adjustable delay**
 - load / calculate / unload modes
 - 16 bit / 12 bit / 8 bit data precision

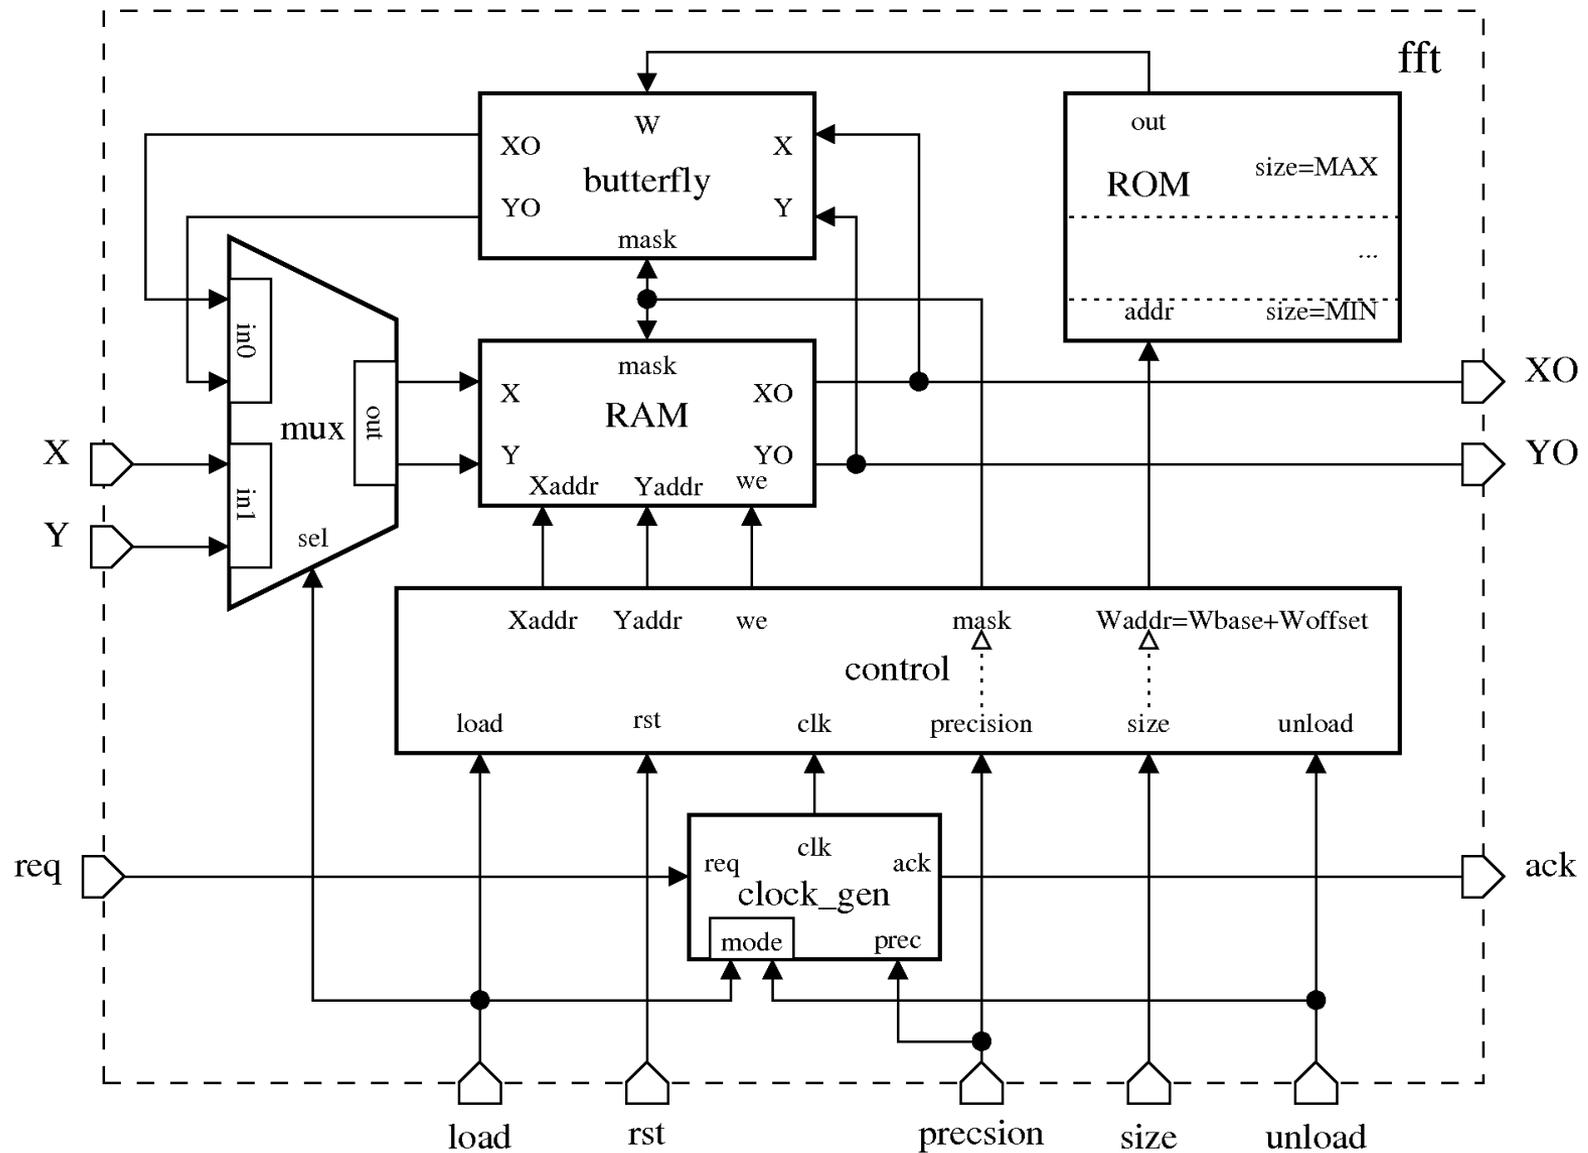
Iterative FFT: Non-reconfigurable



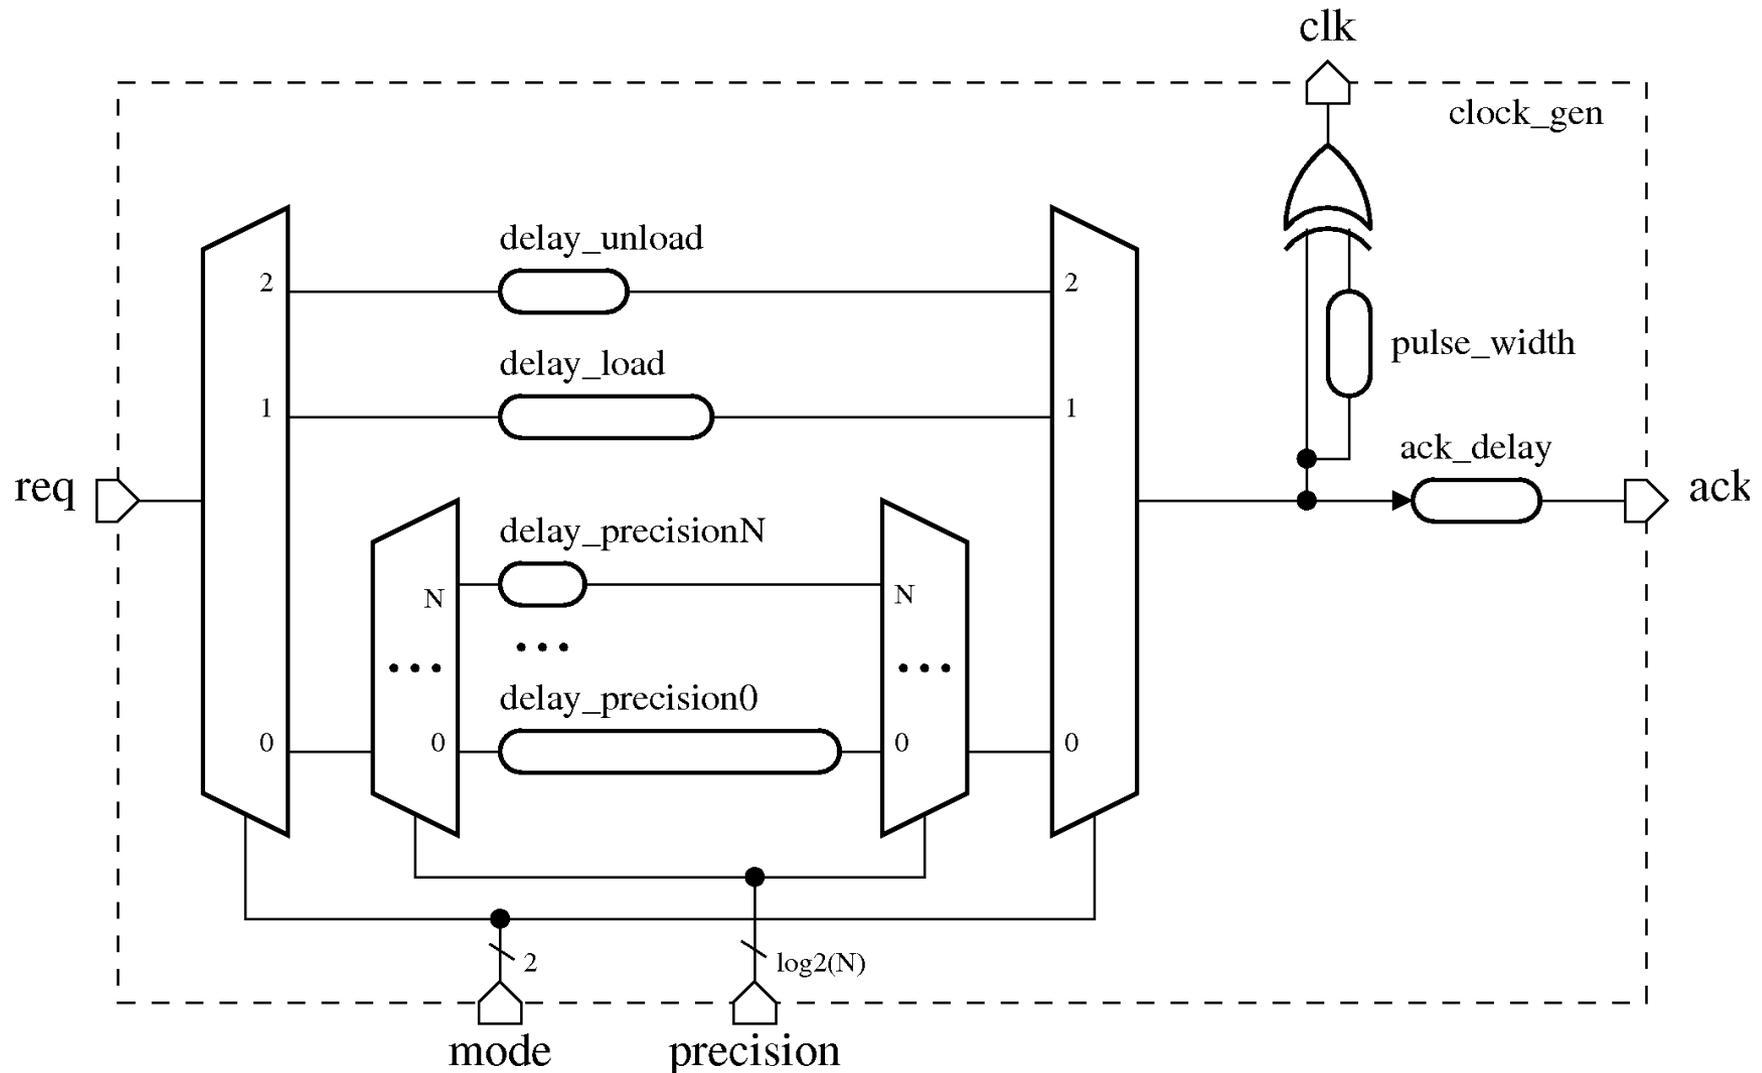
Iterative FFT: Reconfigurable



Iterative FFT: Asynchronous



Iterative FFT: Clock Generator

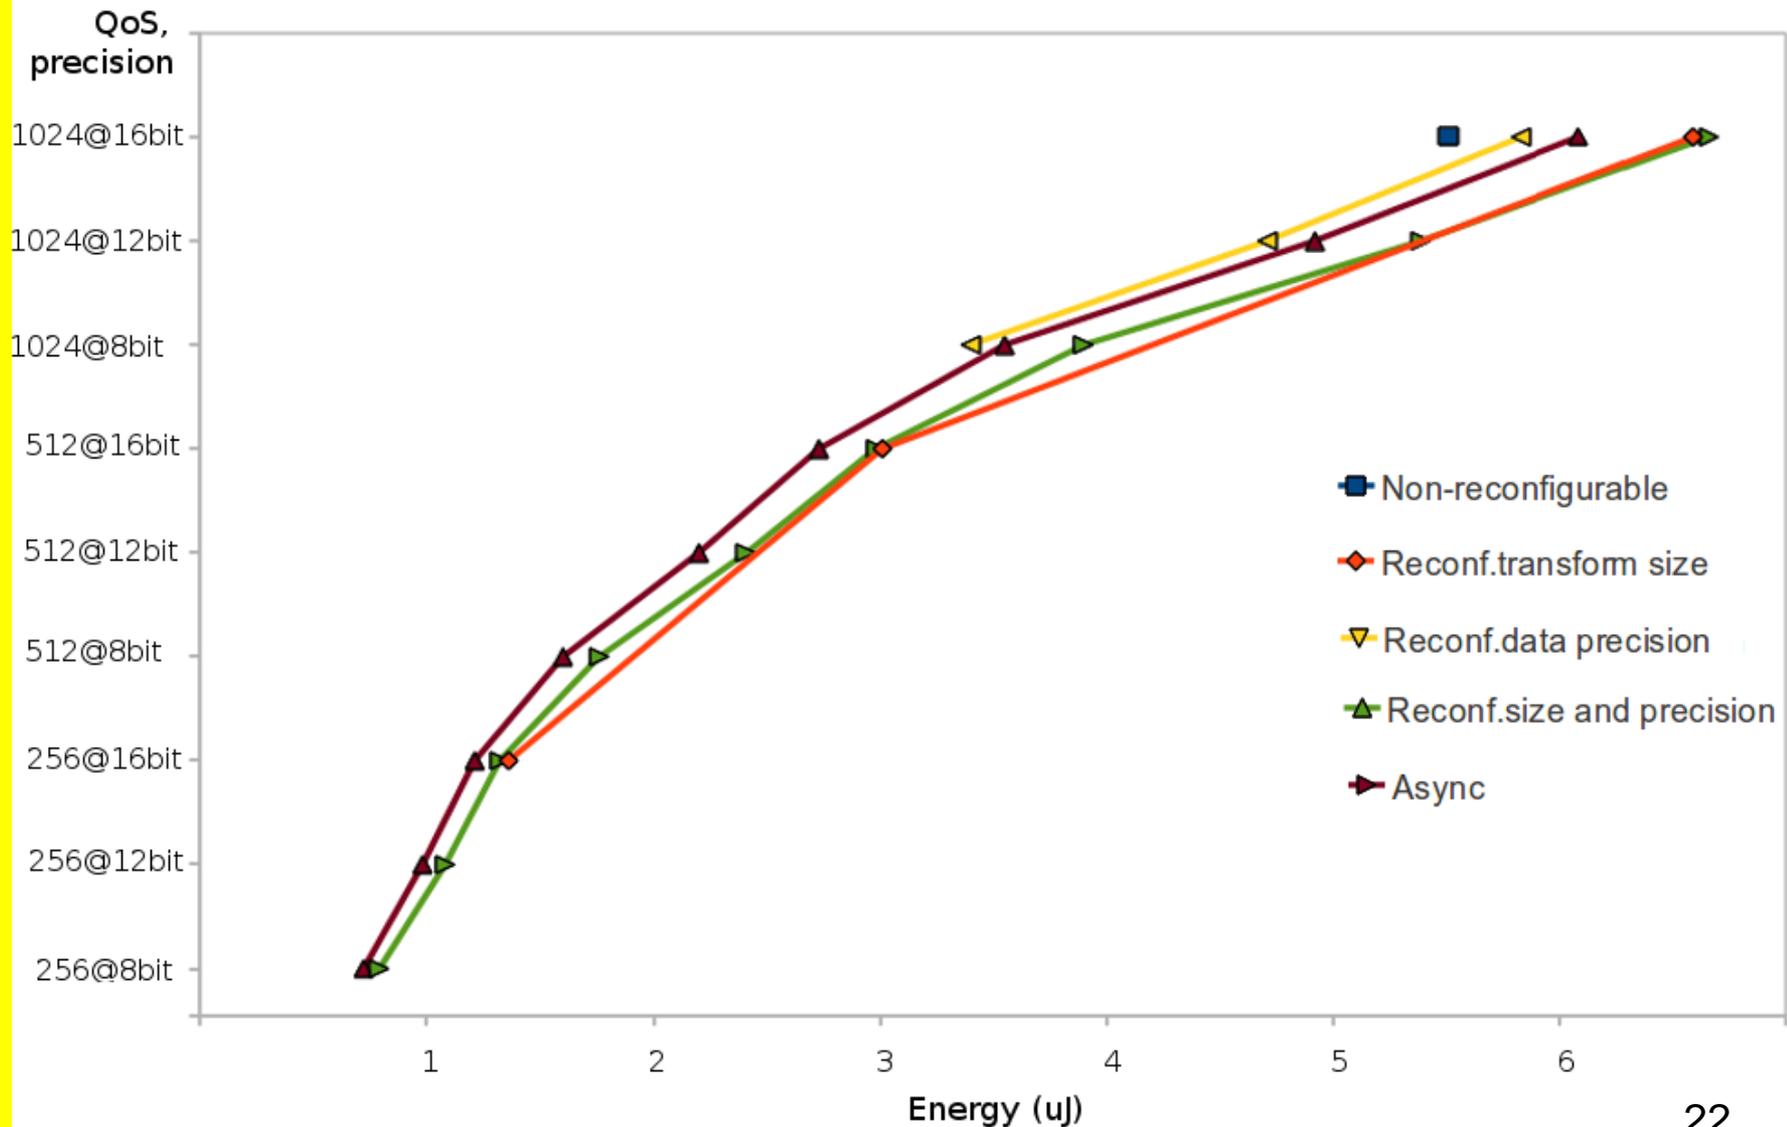


Iterative FFT: Simulation Results

	Area (cells)	Energy consumption (μJ)								
		1024 points			512 points			256 points		
		16bit	12bit	8bit	16bit	12bit	8bit	16bit	12bit	8bit
A	914K	5.51	-	-	-	-	-	-	-	-
B	920K	6.59	-	-	3.01	-	-	1.36	-	-
C	925K	5.83	4.71	3.40	-	-	-	-	-	-
D	936K	6.66	5.39	3.90	2.98	2.41	1.76	1.32	1.08	0.79
E	942K	6.09	4.92	3.55	2.73	2.20	1.60	1.21	0.98	0.72

- A - non-reconfigurable
- B - reconfigurable transform size
- C - reconfigurable data precision
- D - reconfigurable size and precision
- E - asynchronous reconfigurable

Iterative FFT: Quality of Service



Towards Power-Adaptive Systems

- **Truly energy-modulated design must be power-adaptive**
- **Systems that are power adaptive are more resourceful and more resilient**
- **Power-adaptive systems can work in a broad range of power levels**
- **How to design such systems?**
- **Let's first consider the difference between battery-powered and harvester-powered system designs ...**

Portable Power Supplies

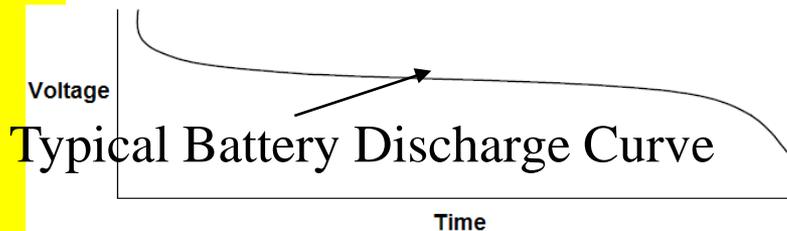
For mobile computing applications the choices of power supply are either batteries or emerging energy-harvester supplies.

Battery:

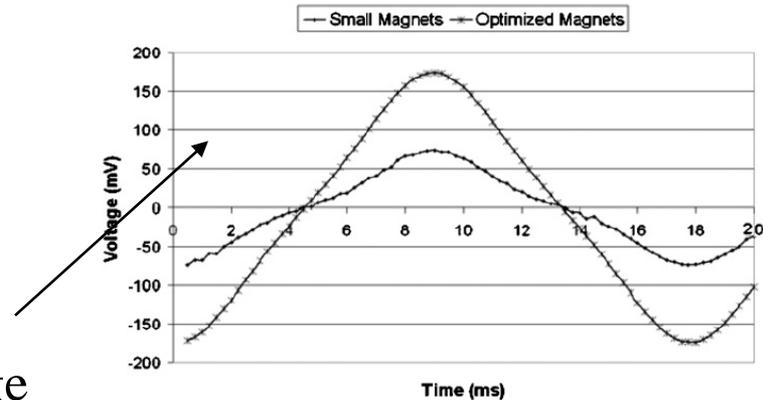
- Can supply finite energy (E) – depends on the battery capacity.
- The available power (P) can be very large.

Energy-Harvester:

- Can supply infinite energy (E).
- The rate of energy production ($dE/dt = P$) is variable and can be small.



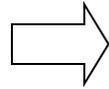
A Micro-Electromagnetic vibration harvester output voltage



S P Beeby et al., 2007, “A micro electromagnetic generator for vibration energy harvesting”, J. Micromech. Microeng. 17 (2007) 1257–1265.

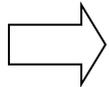
Circuit Designer Choices (1)

Battery
Supply



- Determine from T0 the required power consumption P_0 .
- Design the circuit for constant P_0 consumption
→ constant V_0 supply → constant f_0 performance
(or apply DVS and DVFS to maximise battery life)

Energy-
Harvester
Supply



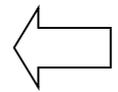
- Design the circuit for constant P_{min} consumption
→ constant V_{min} supply → constant f_{min} performance.
- OR
- Track available power $P_{average}$ → change circuit consumption/performance in real-time → $f_{average} > f_{min}$.

Circuit Designer Choices (2)

To maximise a circuit's power utilization of a variable power output source:

- increase voltage supply of the circuit to the energy-optimum value (variable voltage).
- switch on/off parts of the circuit (constant voltage).

Real-time

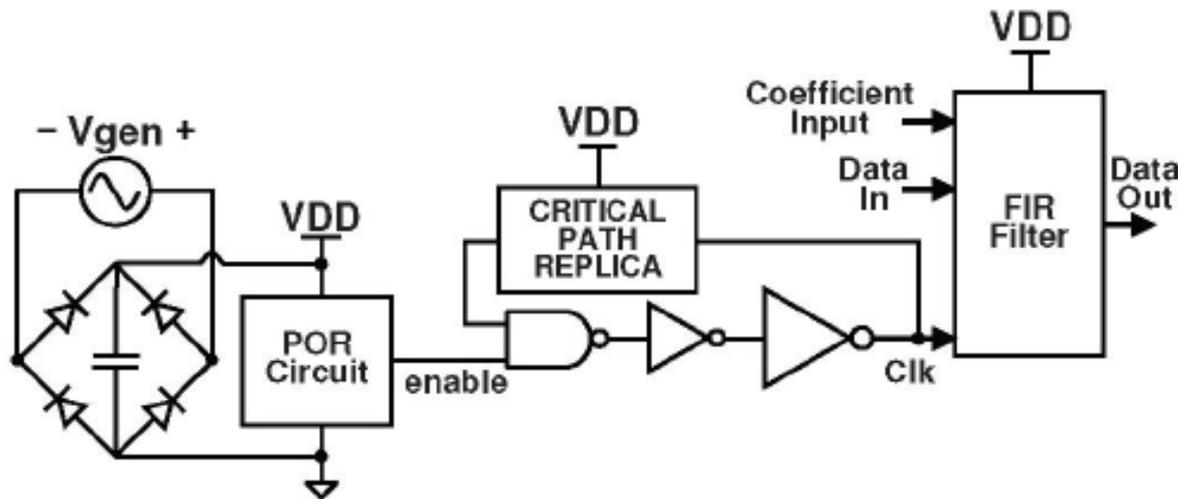


- For both cases special controller circuits have to be developed.
- For the first case (variable voltage) self-timed circuits have an advantage → no additional circuit required to change the operating frequency.

AC supplied self-timed circuits have been demonstrated in practice.

For every power supply cycle: wake up the circuit, perform computation and shut down the circuit – hence, power-on reset needed.

AC-powered self-timed circuit

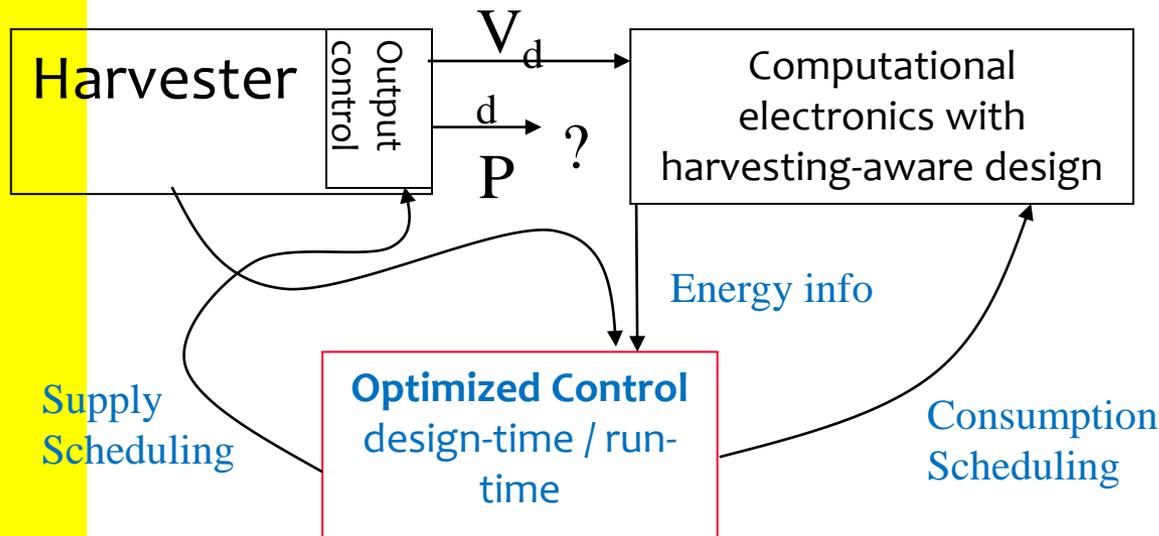


Fast Power-on Reset (4.1nW),
3T DRAM to keep state across supply
cycles,
135K transistors in 180nm CMOS
Can supply 250KHz on all process corners
for $\leq 50^{\circ}\text{C}$

J Wenck, R Amirtharajah, J Collier and J
Siebert, 2007, "AC Power Supply
Circuits for Energy Harvesting", 2007 IEEE
Symposium on VLSI Circuits, 92-93.

**Problems: critical path replica may not scale well with the
computational load (cf. SRAM delay matching problems – following
slides)**

Power-adaptive Computing (Holistic view)



Optimization

$$\frac{E_C}{E_S} = \text{Max}$$

**Useful energy consumption is maximized for a given amount of energy produced, or
Energy supplied is minimized for given amount of energy consumed usefully (to carry out specified/required computation)**

Power-Adaptive System Design

Adaptation levels:

- **Cell and component level**
 - Resilience to Vdd variations (e.g. robust synchronisation, self-timed logic and completion detection)
 - Leakage control mechanisms (e.g. body biasing)
- **Circuit level**
 - Clock/power gating, DVF scaling
 - Reconfiguration of logic (turning on-off parts, concurrency control, completion detection on-off)
- **System level (power sensing and control of power supply and consumption chains)**
 - Optimal control of Vdd for minimum energy per operation
 - Control of computation load to fit the power profile or optimise for average power

Power-Adaptive System Design

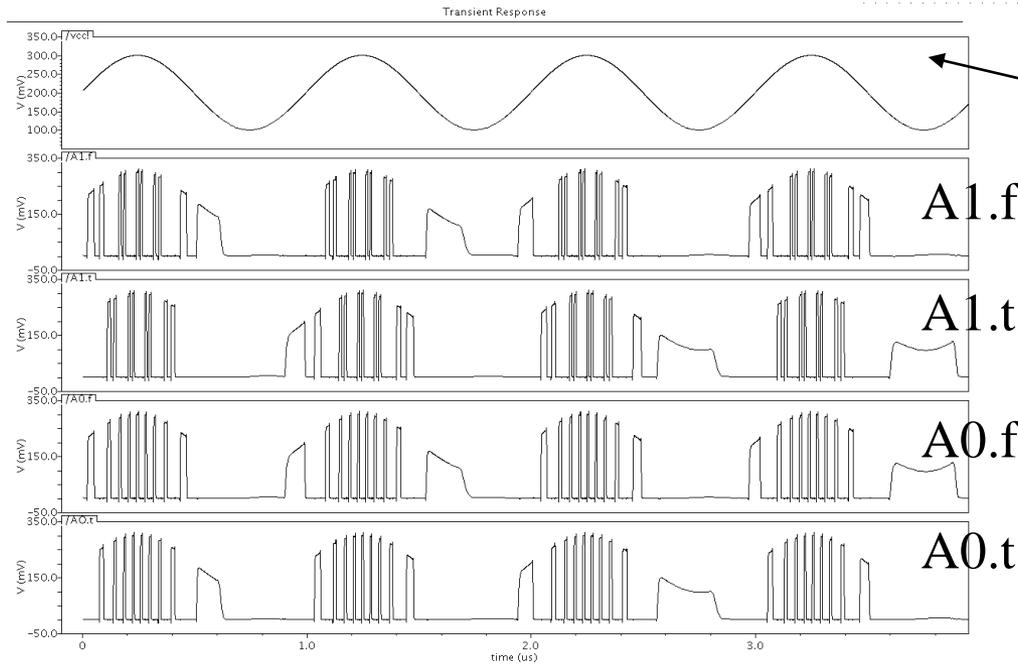
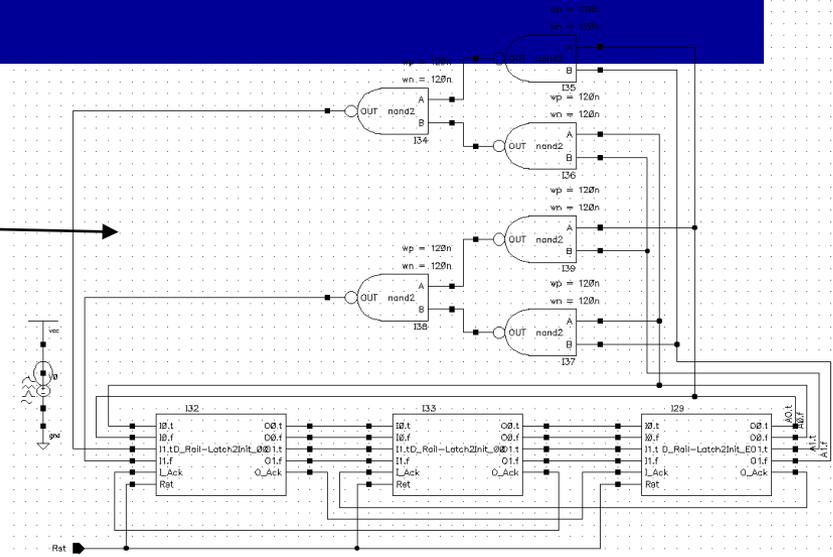
Adaptation levels:

- Cell and component level
 - Resilience to Vdd variations (e.g. robust synchronisation, self-timed logic and completion detection)
 - Leakage control mechanisms (e.g. body biasing)
- Circuit level
 - Clock/power gating, DVF scaling
 - Reconfiguration of logic (turning on-off parts, concurrency control, completion detection on-off)
- System level (power sensing and control of power supply and consumption chains)

Asynchronous (self-timed) design principles improve effectiveness and efficiency of both sensing and control in adaptation process

Closer look at AC-powered self-timed logic

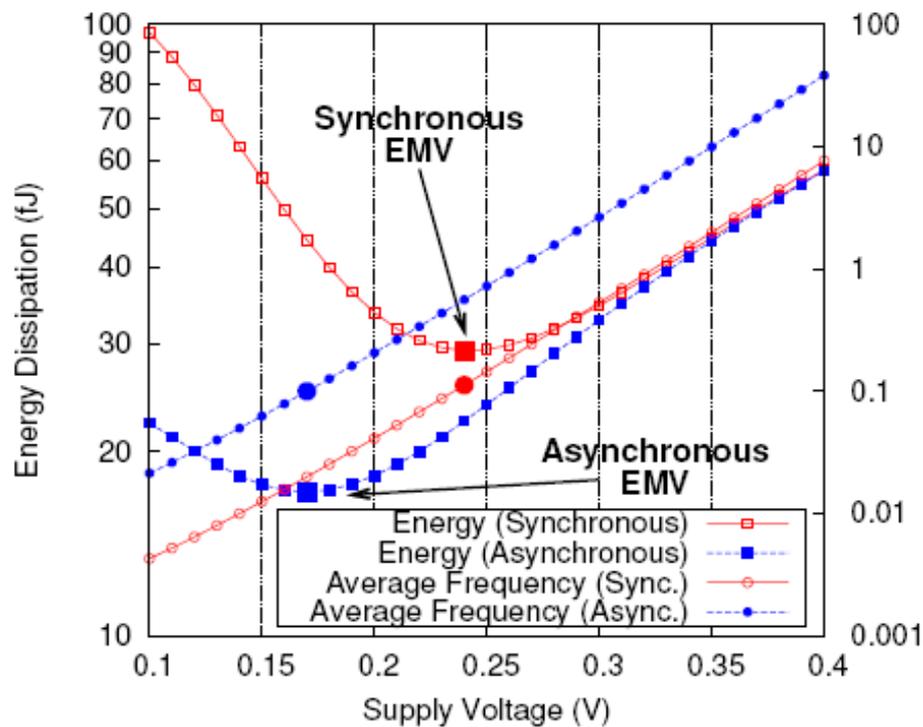
2-bit Sequential Dual-rail Asynchronous Counter



Supply: AC
200mV±100mV
Frequency: 1Mhz

Self-timed logic with completion detection is robust to power supply variations

Synchronous vs Self-Timed Design (in terms of energy efficiency)



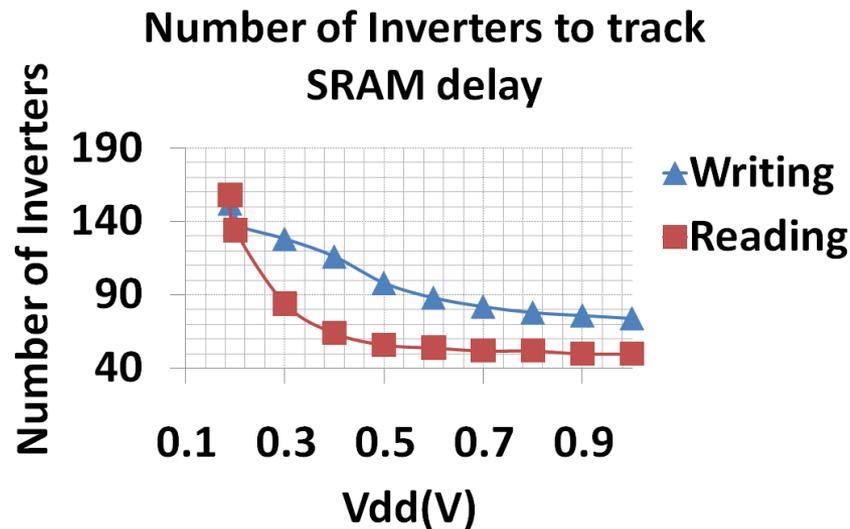
Average Operating Frequency (MHz)

Asynchronous (self-timed) logic can provide completion detection and thus reduce the interval of leakage to minimum, thereby doing nothing well!

Source: Akgun et al, ASYNC'10

Speed-independent SRAM

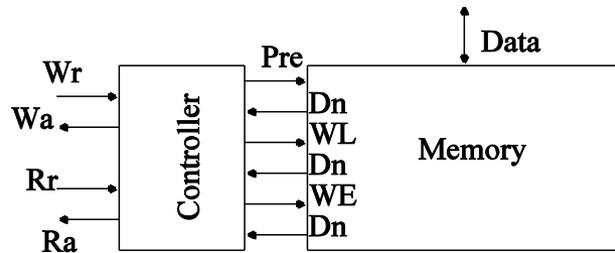
- Mismatch between delay lines and SRAM memories when reducing Vdd



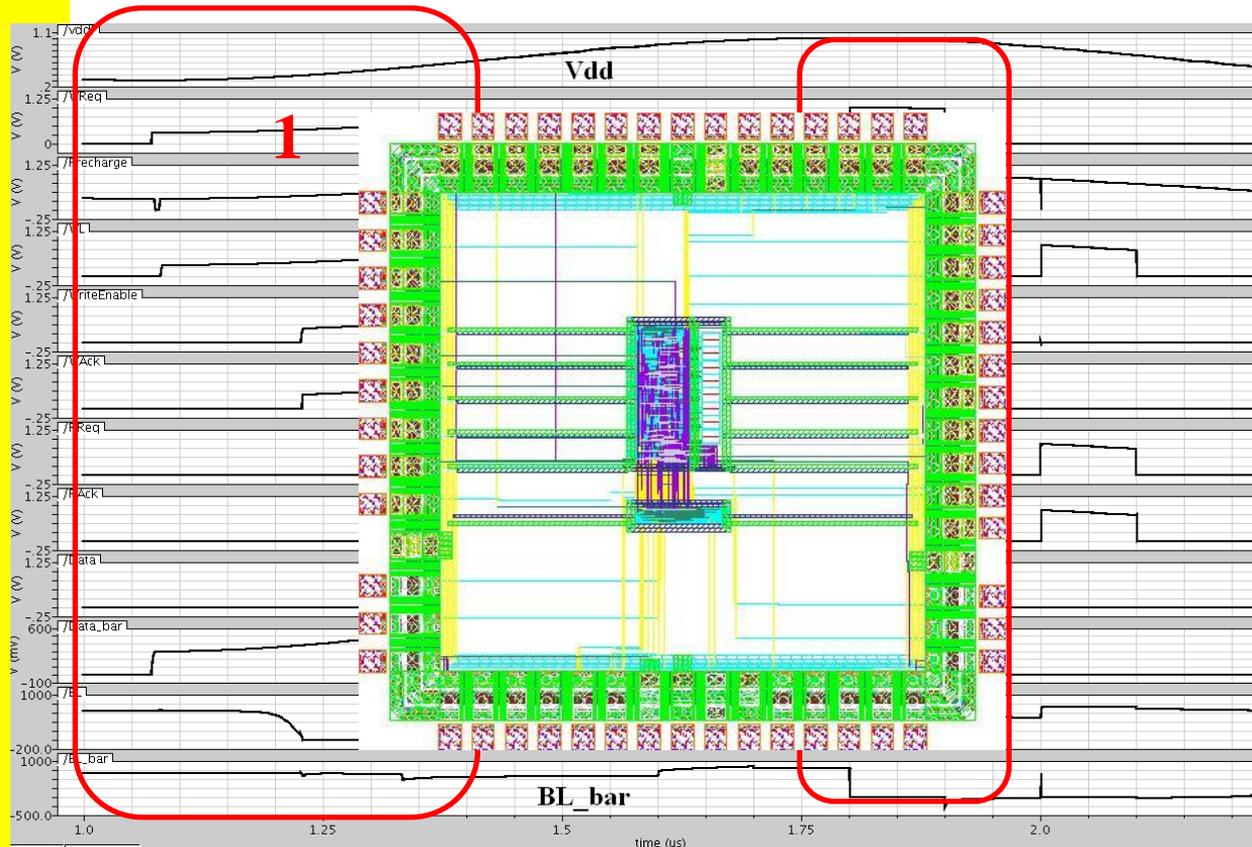
For example, under 1V Vdd, the delay of SRAM reading is equal to 50 inverters and under 190mV, the delay is equal to 158 inverters

- The problem has been well known so far
- Existing solutions:
 - Different delay lines in different range of Vdd
 - Duplicating a column of SRAM to be a delay line to bundle the whole SRAM
- The solutions require:
 - voltage references
 - DC-DC adaptor
- **Completion detection needed?!**

Speed-independent SRAM



The SI controller uses **completion detection in SRAM** and handshake protocols to manage pre-charge, WL and WE in the SRAM banks



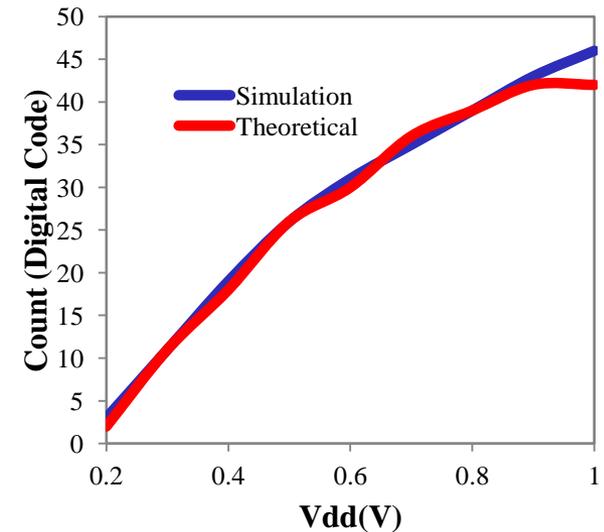
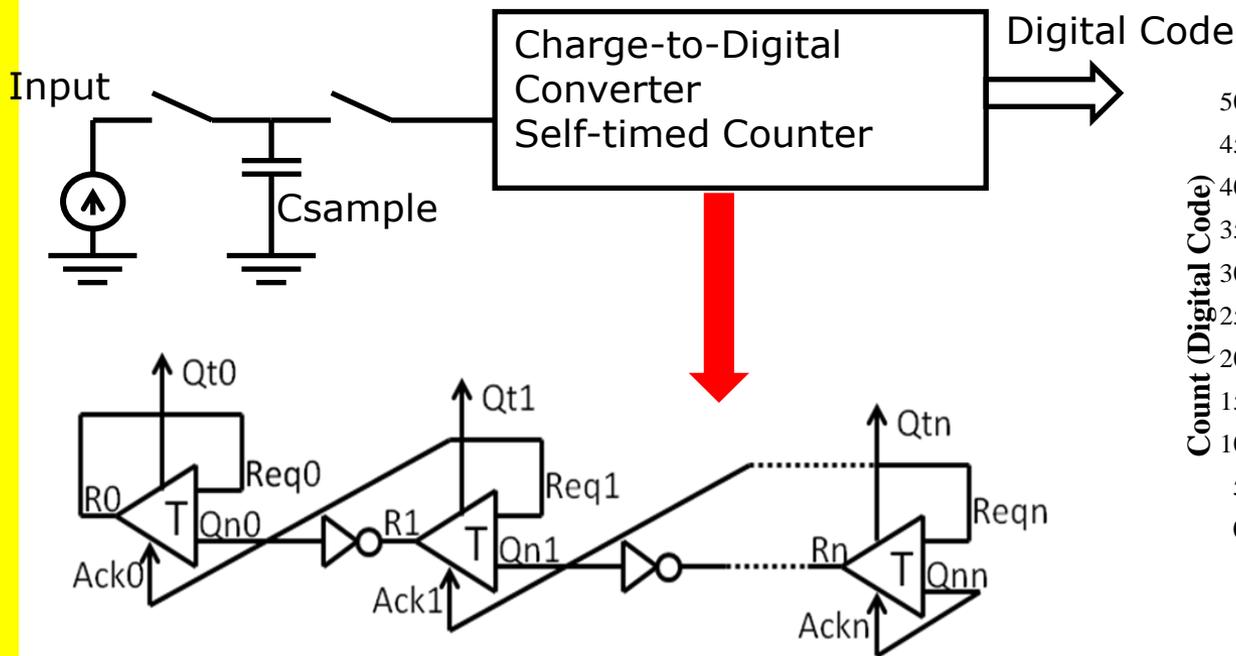
Can work smoothly under variable Vdd.

For example, the first writing works in low Vdd, it takes long time, and the second writing works in high Vdd, works faster.

A. Baz et.al. PATMOS 2010

Voltage sensors for power-adaptive computing

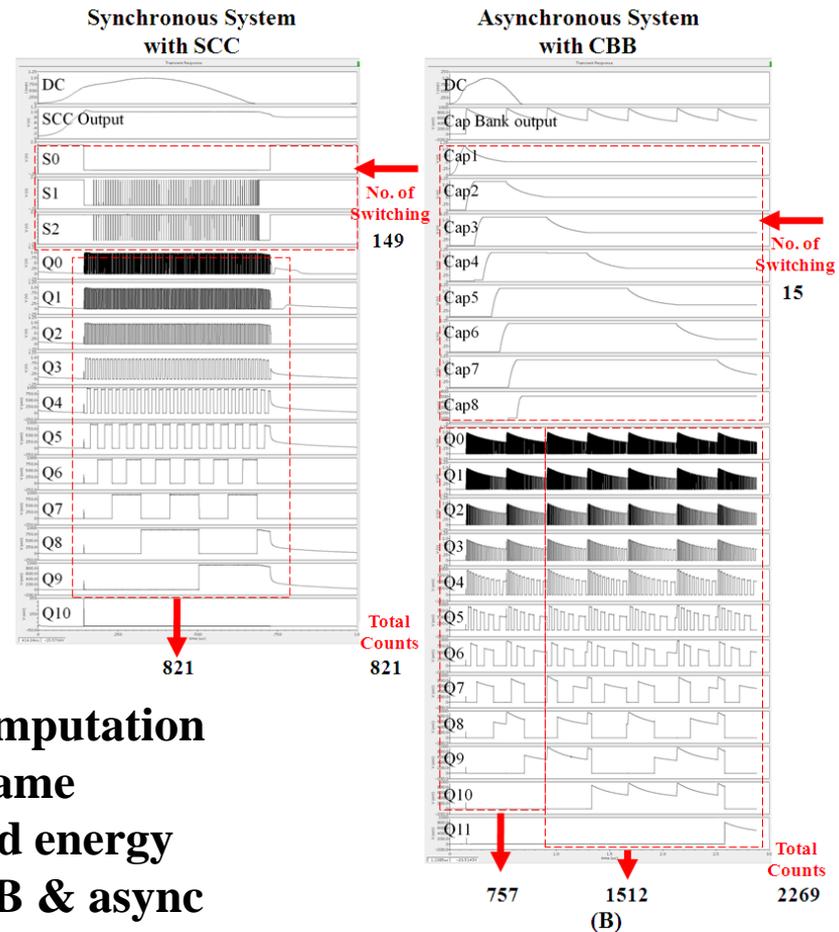
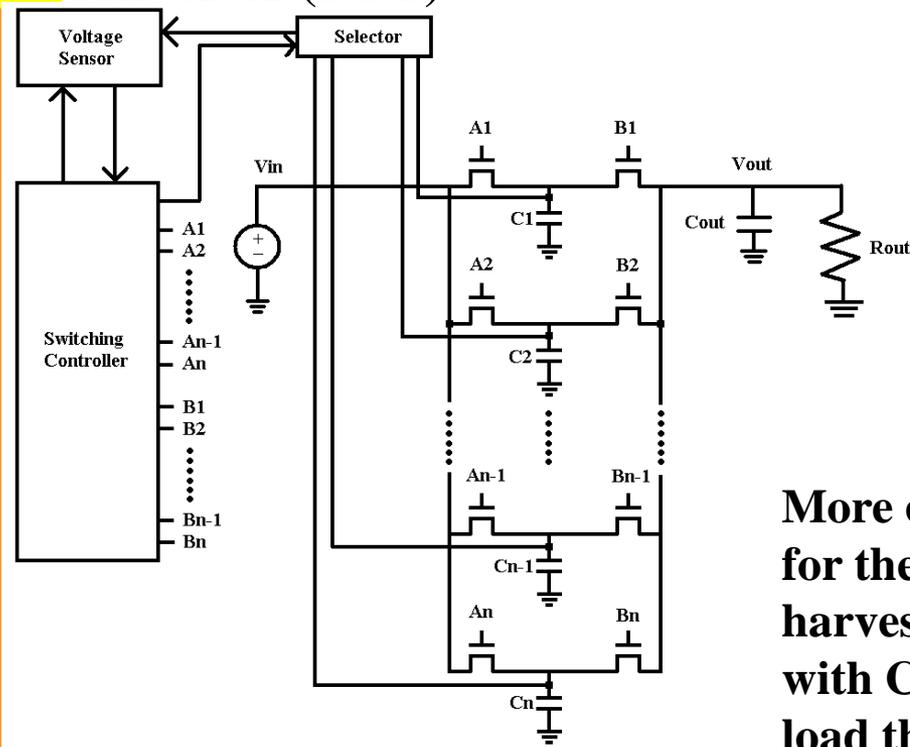
Voltage sensors which can work under variable Vdd without or with minimal reference requirements are indispensable for control and optimization ...



*UK Patent application 1005372.6
(30.03.10), Newcastle University*

Novel power electronics for power-proportional logic

Power control based on capacitor bank blocks (CBBs) instead of DC-DC based on switched capacitor converter (SCC)



More computation for the same harvested energy with CBB & async load than SCC & sync load

X. Zhang et.al. ASYNC 2011

Conclusions

- Energy-harvesting changes the dynamic balance between supply and consumption – supply adds operational constraints in real-time
- Adaptation to power changes should be at all levels of abstraction, from logic cells to systems
- Asynchronous (self-timed) techniques support more effective adaptation to V_{dd} changes via natural temporal robustness; they also offer better energy proportionality
- Good energy characterisation of loads (logic, memory, i/o, RF) is essential for high-quality adaptation
- More theory, models and algorithms are needed for handling the problem of power-adaptation in run-time
 - Good characterisation of systems functions, s/w and h/w tasks in terms of duty cycling (periodic, sporadic, bursty)
 - Similar characterisation of energy sources (p,s,b)
 - Online power monitoring is essential to mitigate uncertainty at design time

Acknowledgements

Members of “Microelectronics Systems Design” research group at Newcastle:

Panagiotis Asimakopoulos, Alex Bystrov, Terrence Mak, David Kinniment, Andrey Mokhov, Delong Shang, Danil Sokolov, Fei Xia, Reza Ramezani, Zhou Yu, Abdullah Baz, Xuefu Zhang, involved in the Power-Adaptive Computing research

EPSRC funding (EP/G066728) – Project “Next Generation Energy-Harvesting Electronics: A Holistic Approach”, involving Universities of Southampton, Bristol, Newcastle and Imperial College

EPSRC ICT Dream Fellowship

Aims of Dream Fellowship:

- develop **creative processes** in the Newcastle research environment and outside, in the area of **energy-modulated computing**, identify challenges, formulate research problems, leading to new research projects and grants, active networking within academia and with industry
- (more specifically) help creating new methods and tools for resource-driven system design

For this an extensive program of visits and interactions with the leading groups:

- UC Berkeley, Stanford, Michigan, MIT, EPFL, ETH, Bologna, CMU, Intel, ARM ...

Time frame: September 2011-August 2013