Newcastle University

holistic energy harvesting

Energy-Modulated Computing

Alex Yakovlev, School of EECE, Newcastle University

Holistic Project Showcase Event, Imperial College London, 11th Feb 2013

Outline

- Introduction: Energy-uncertainty-QoS interplay
- Energy-modulated computing: approach, principles, system design, examples
- Power-proportional systems
  - Power proportional and power efficient systems
  - Experiment in power-proportionality
  - System Design for energy-harvester power supply
- Power-adaptive systems
  - SRAM, Voltage sensors, Novel power electronics
- Conclusion

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, …**

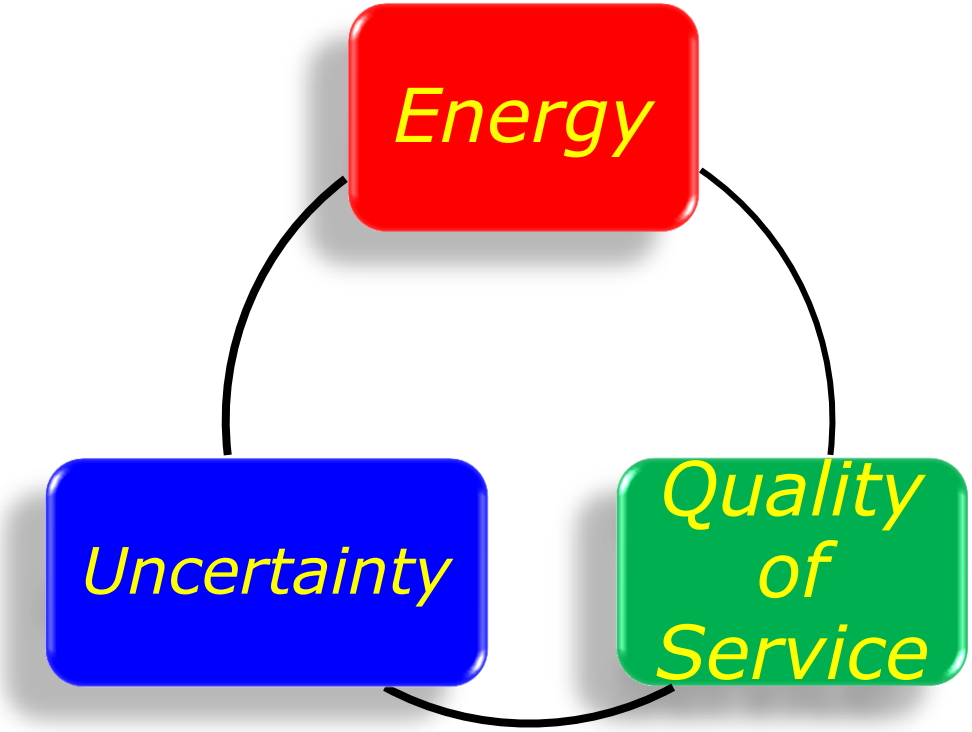


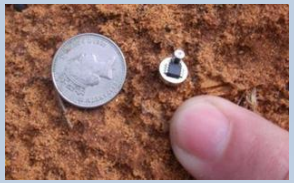

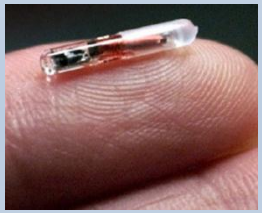
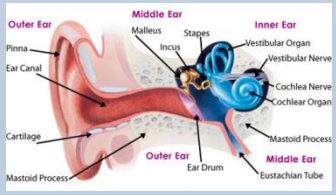


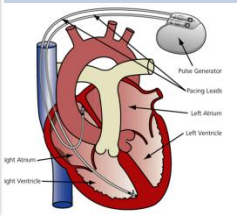

Power/Energy modulation

- The principle of **power/energy-modulated computing** is that the flow of energy entering a computing system determines its computational flow
- It is fundamental for building future systems, particularly for **survivability**
- Any piece of electronics becomes active and performs to a certain level of its delivered **quality in response to some level of energy and power**
- A **quantum of energy** when applied to a computational device can **be converted** into a corresponding **amount of computation activity**
- Depending on their design and implementation systems can produce **meaningful activity at different power levels**
- As **power levels become uncertain** we cannot always guarantee completely certain computational activity

Traditional vs energy-modulated view

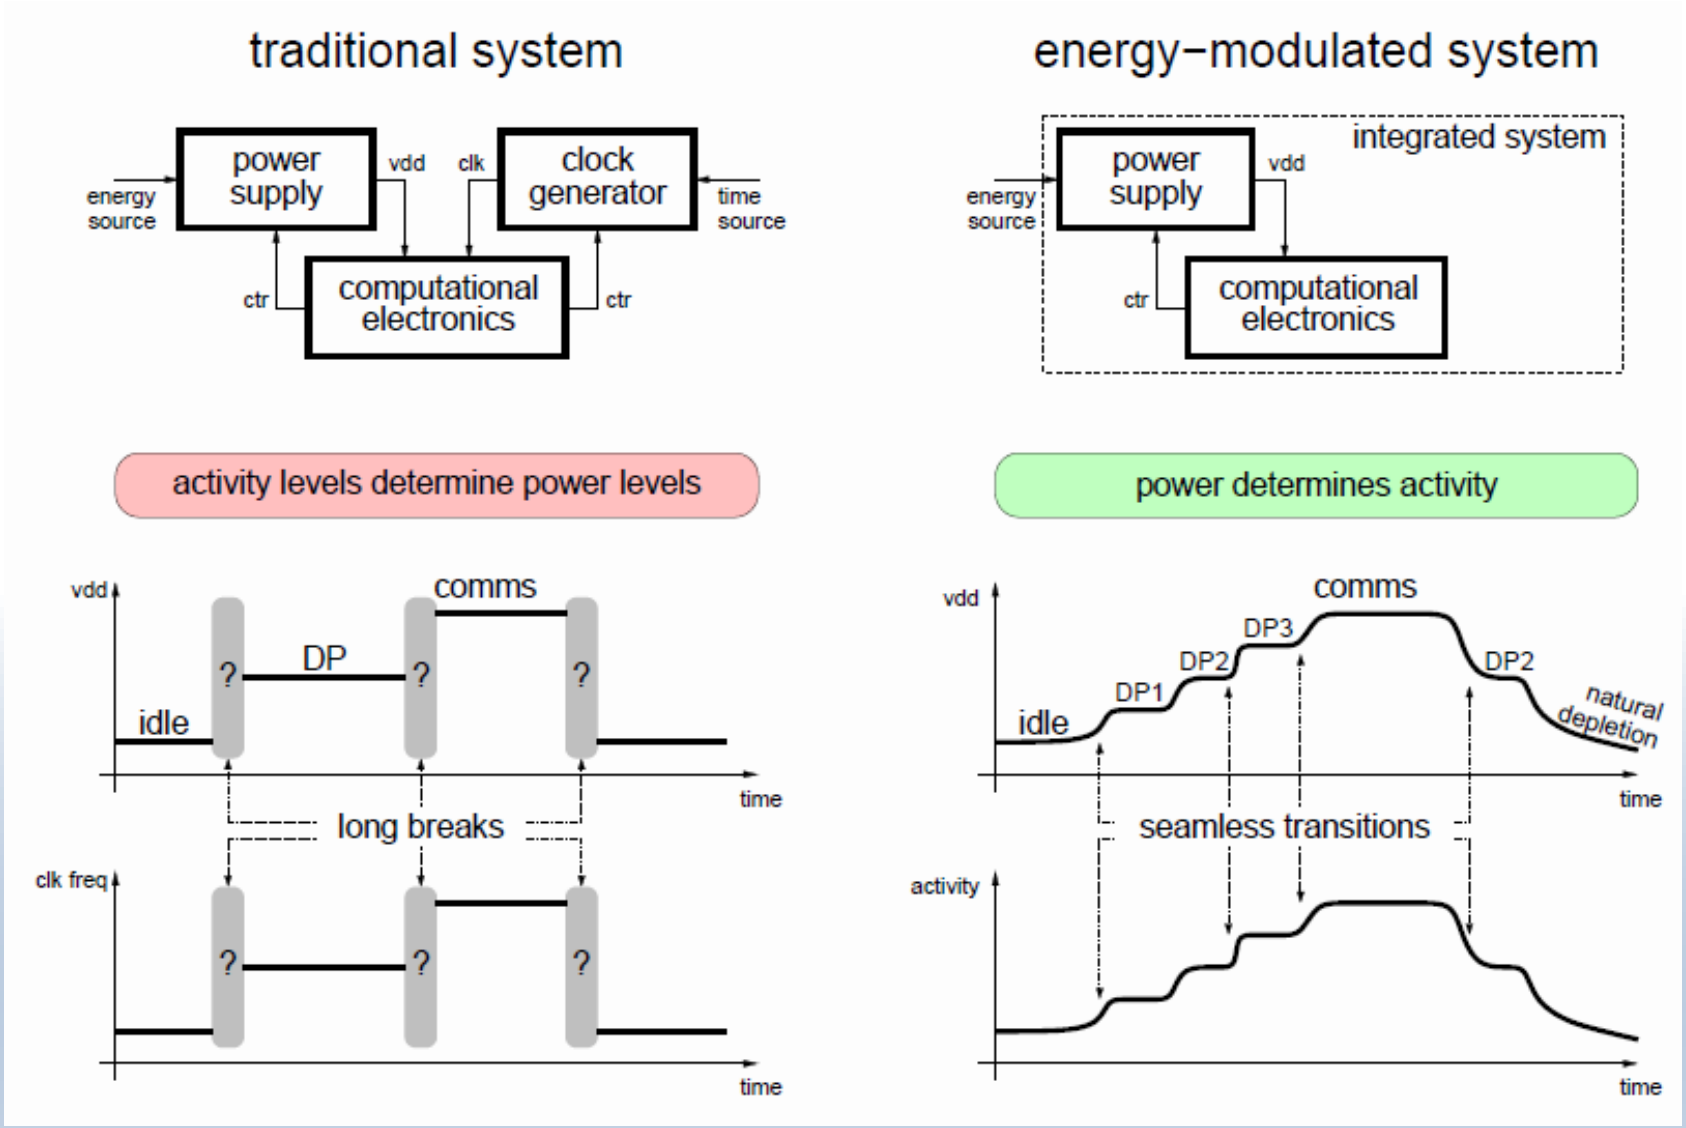

Example 1: condition monitoring

- Energy field (thermal, mechanical vibration, etc)

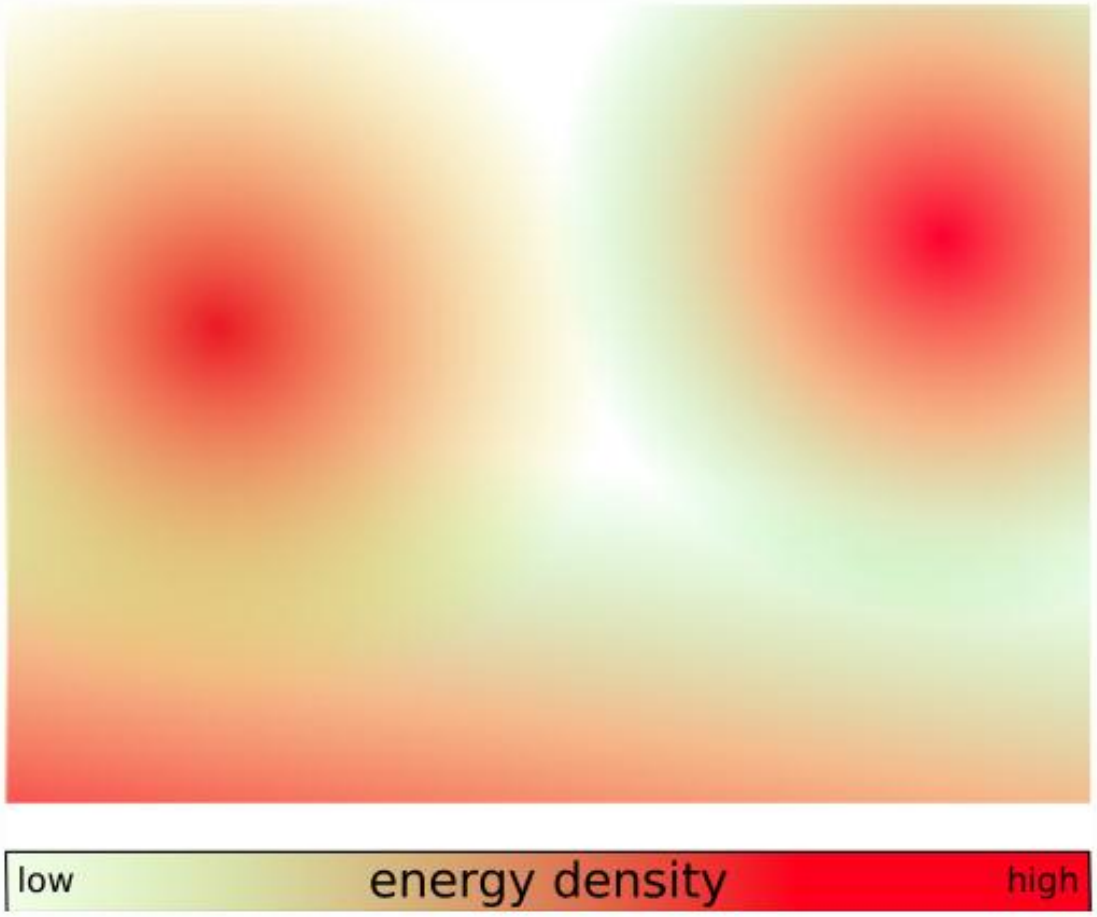

Example 1: condition monitoring

- Network of sensors for spaced and temporal energy mapping

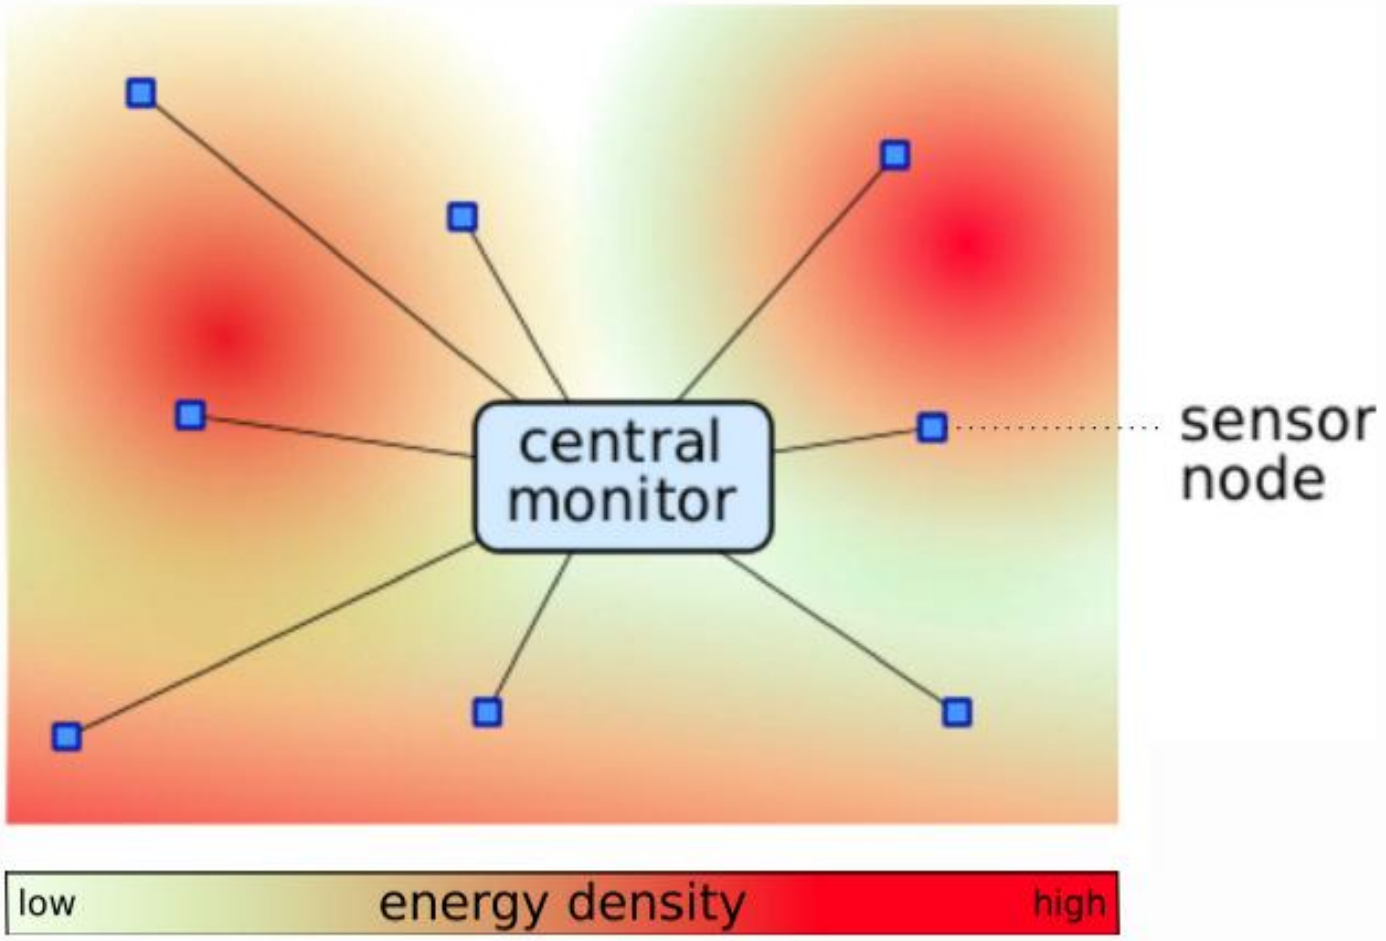

Example 1: condition monitoring

- Sensor node structure

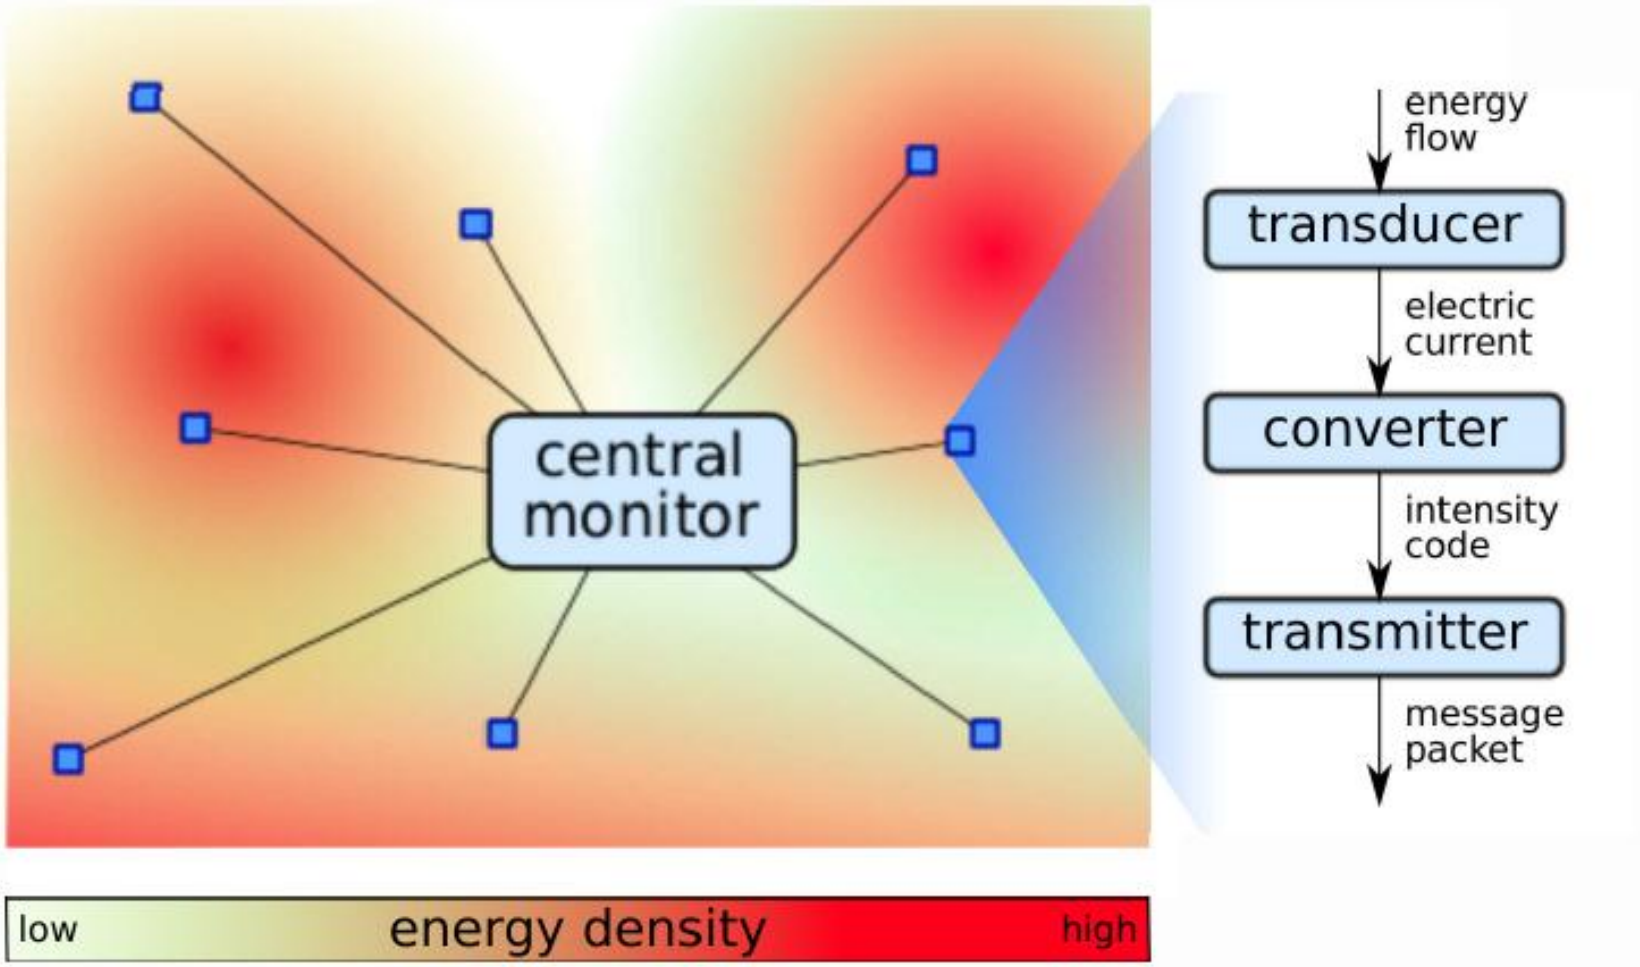

- System architecture

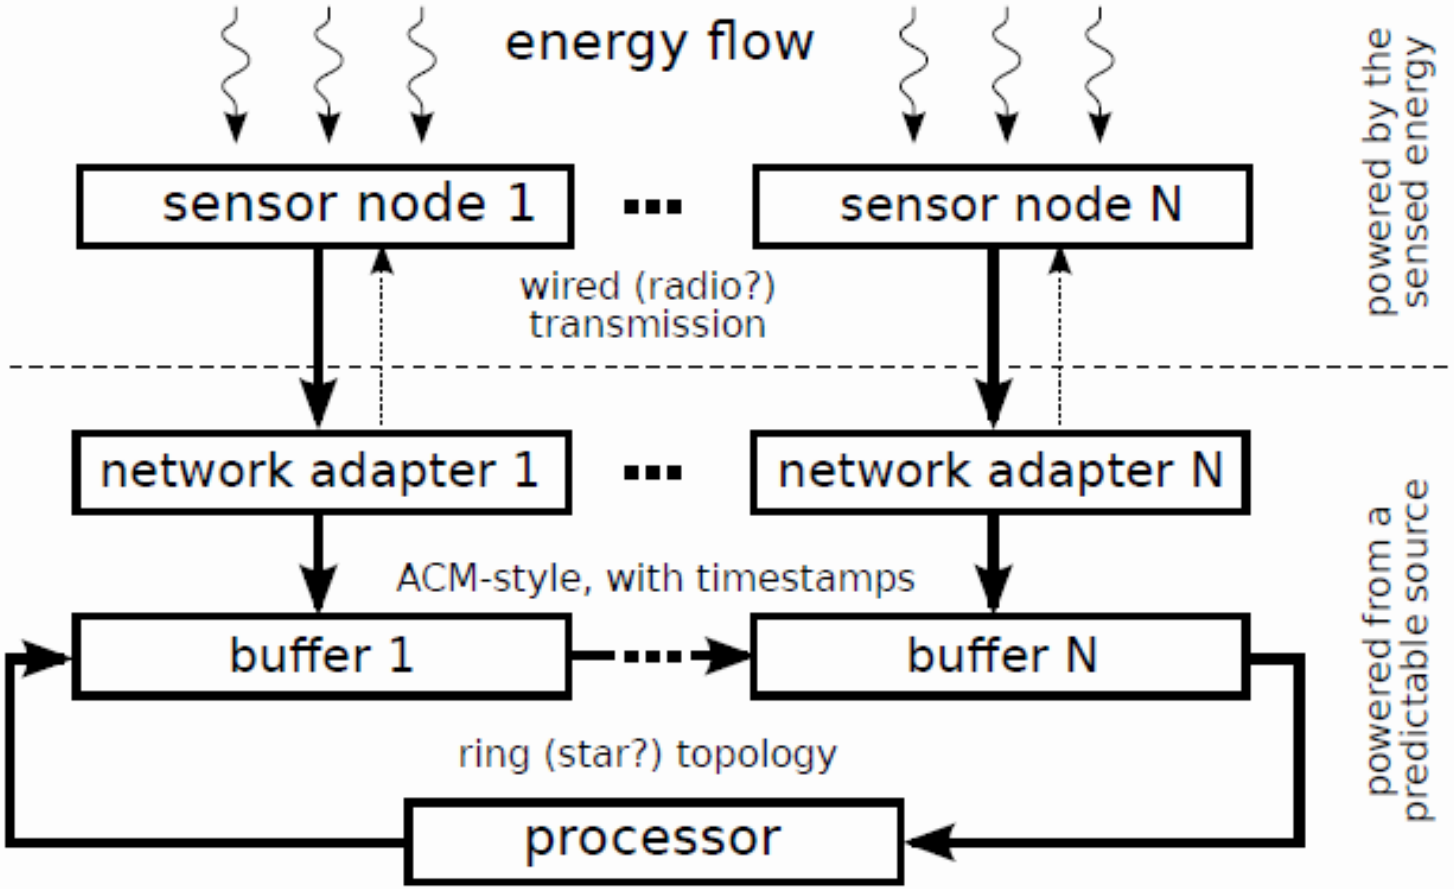

Power-modulation in time

- Localised prediction, from every moment at present
- Power has a certain profile (time trajectory) in the past and uncertain future
- Power-proportional and power-adaptive systems …

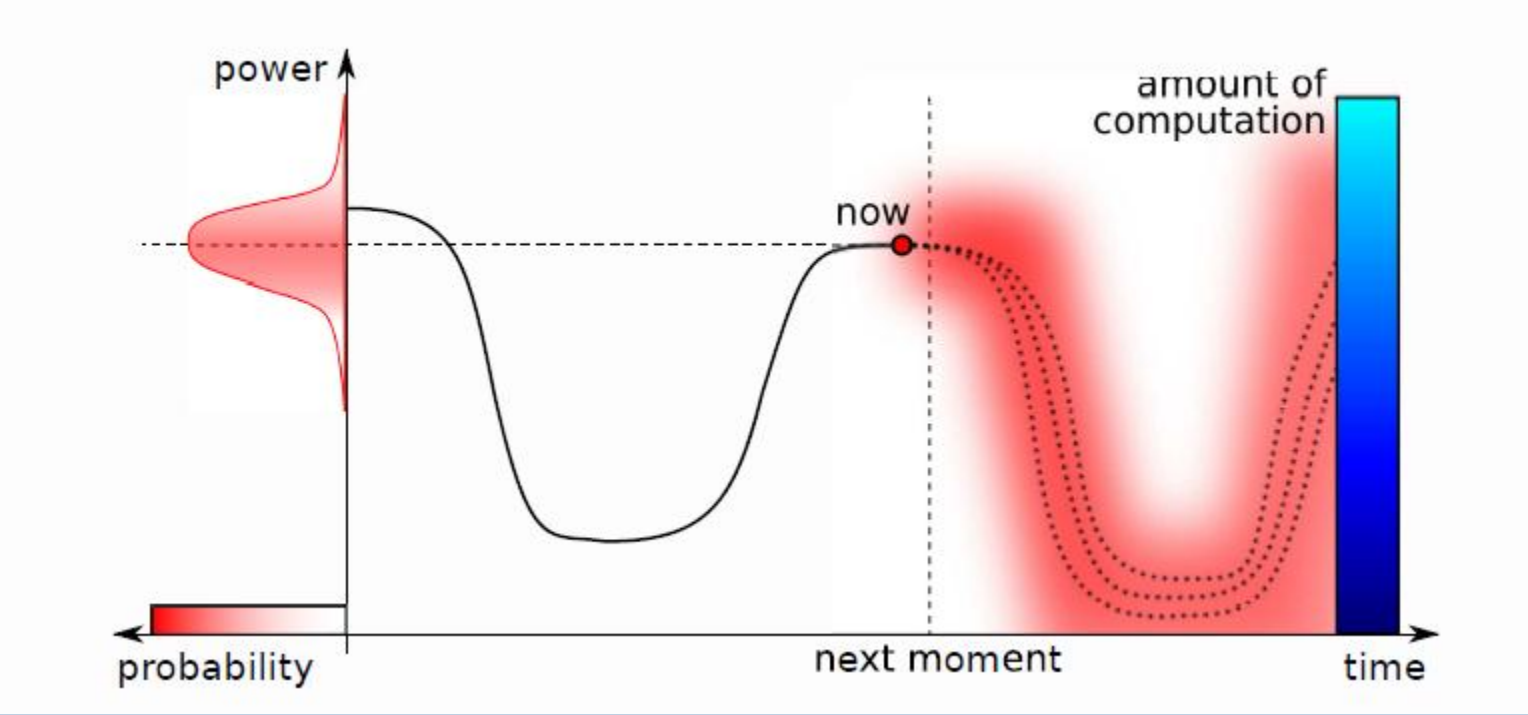

Power proportionality: two views

Energy optimisation for required service demand

Service provision optimisation for constrained power supply

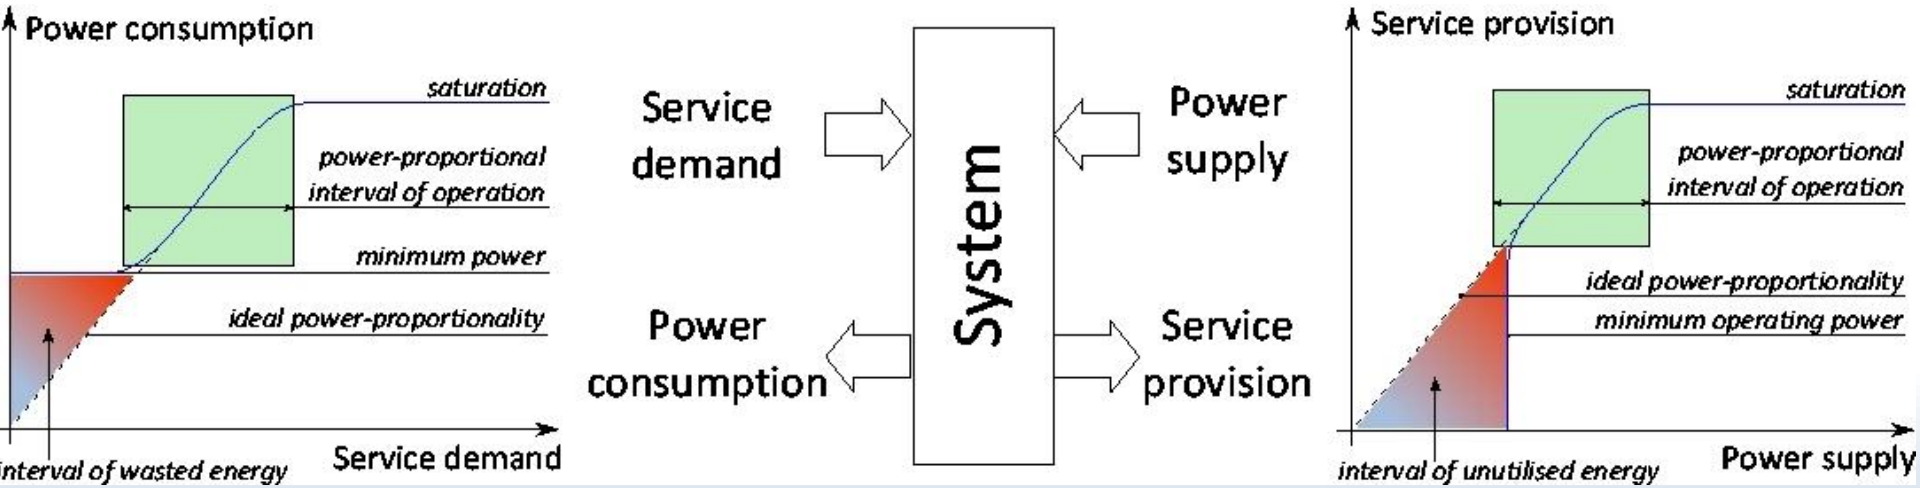

Service-modulated processing

Energy-modulated processing

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

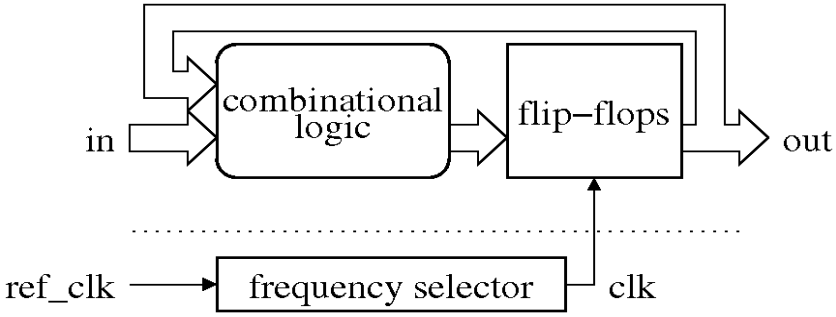

Adaptive frequency scaling

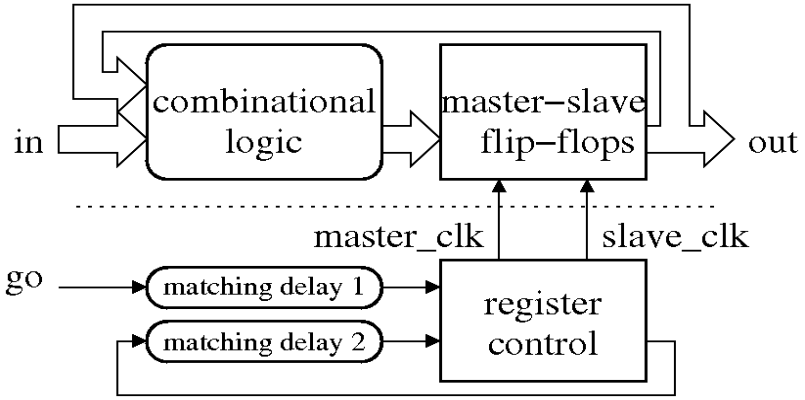

Bundled data

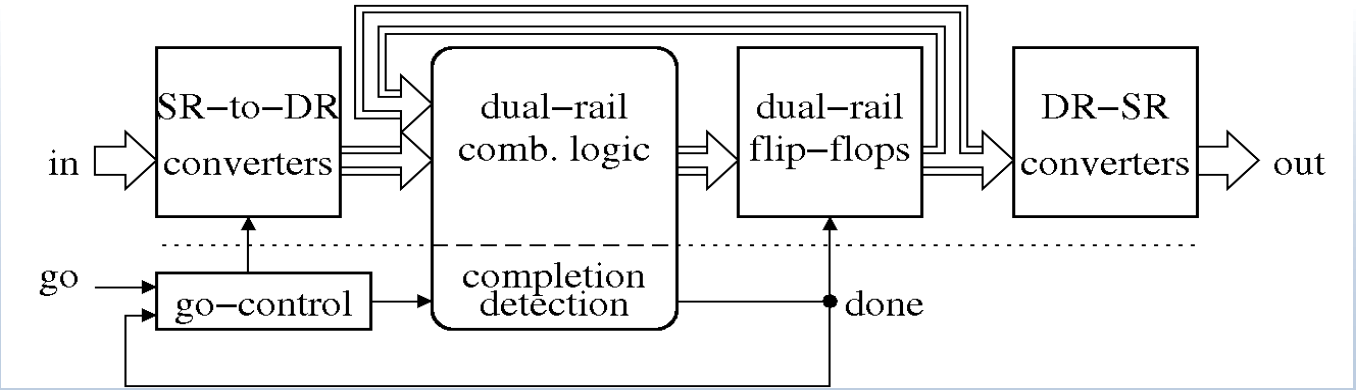

Dual-rail

Multiplier: Quality of Service

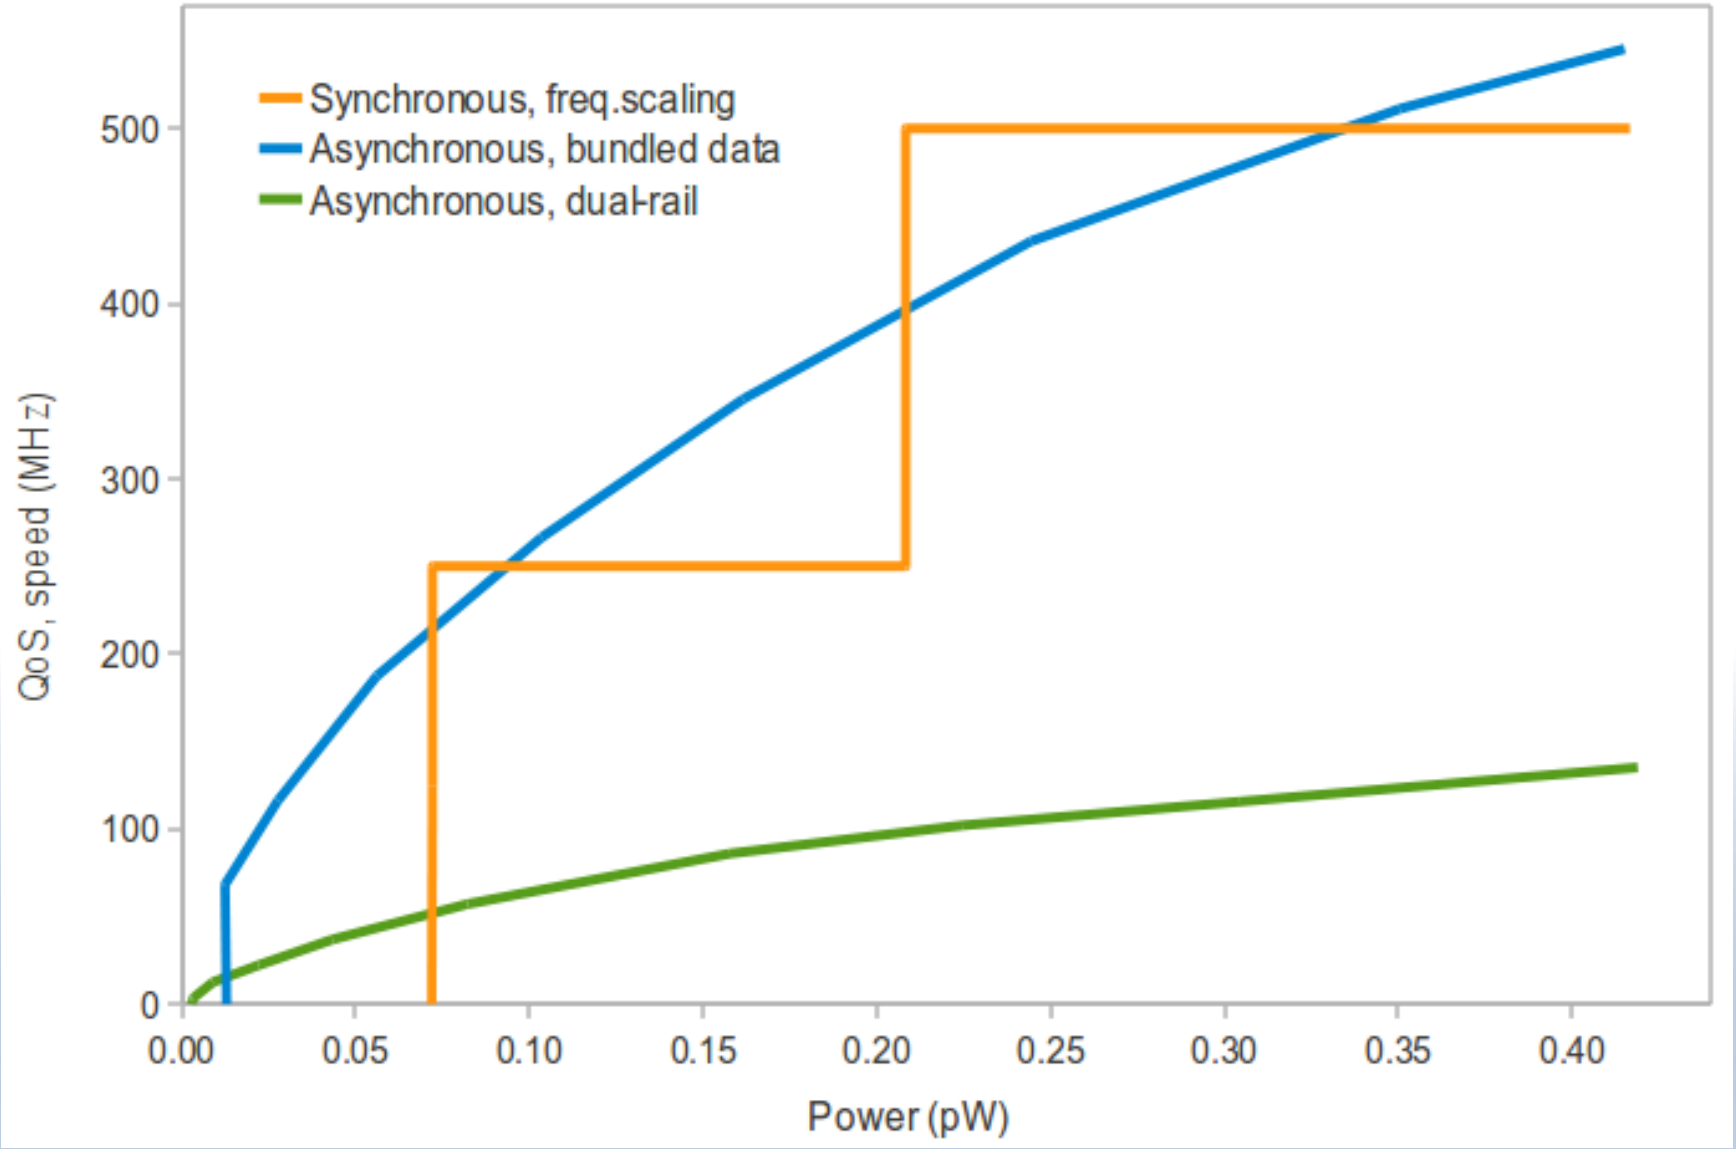

From power-proportional to power-adaptive

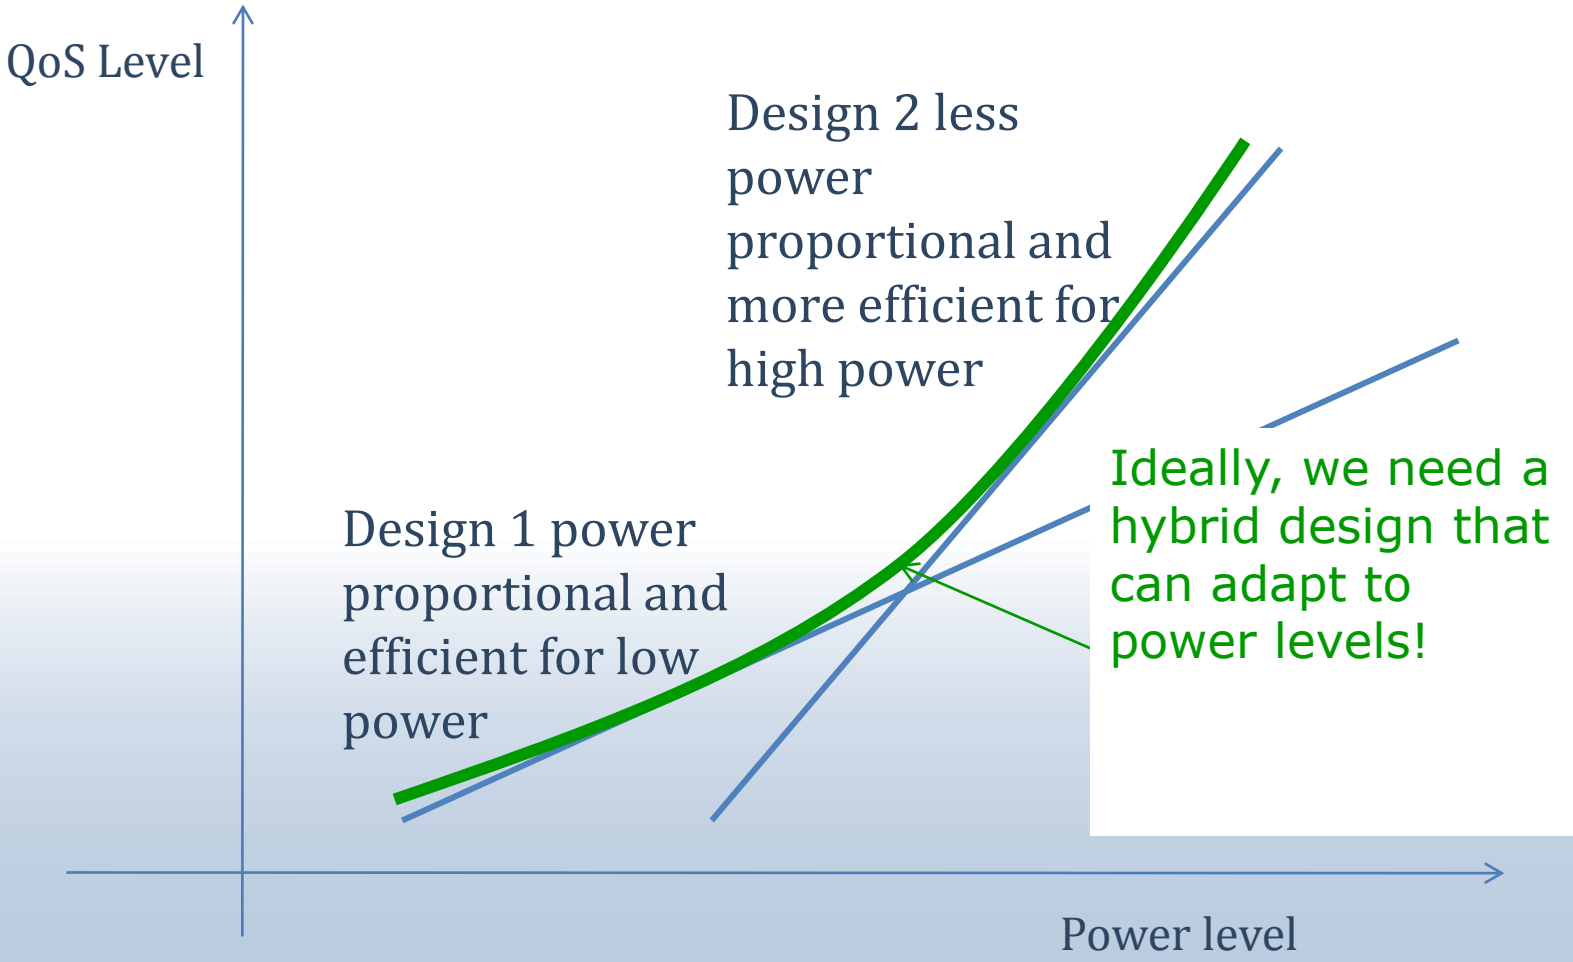

Relationship with uncertainty (e.g. timing variability)

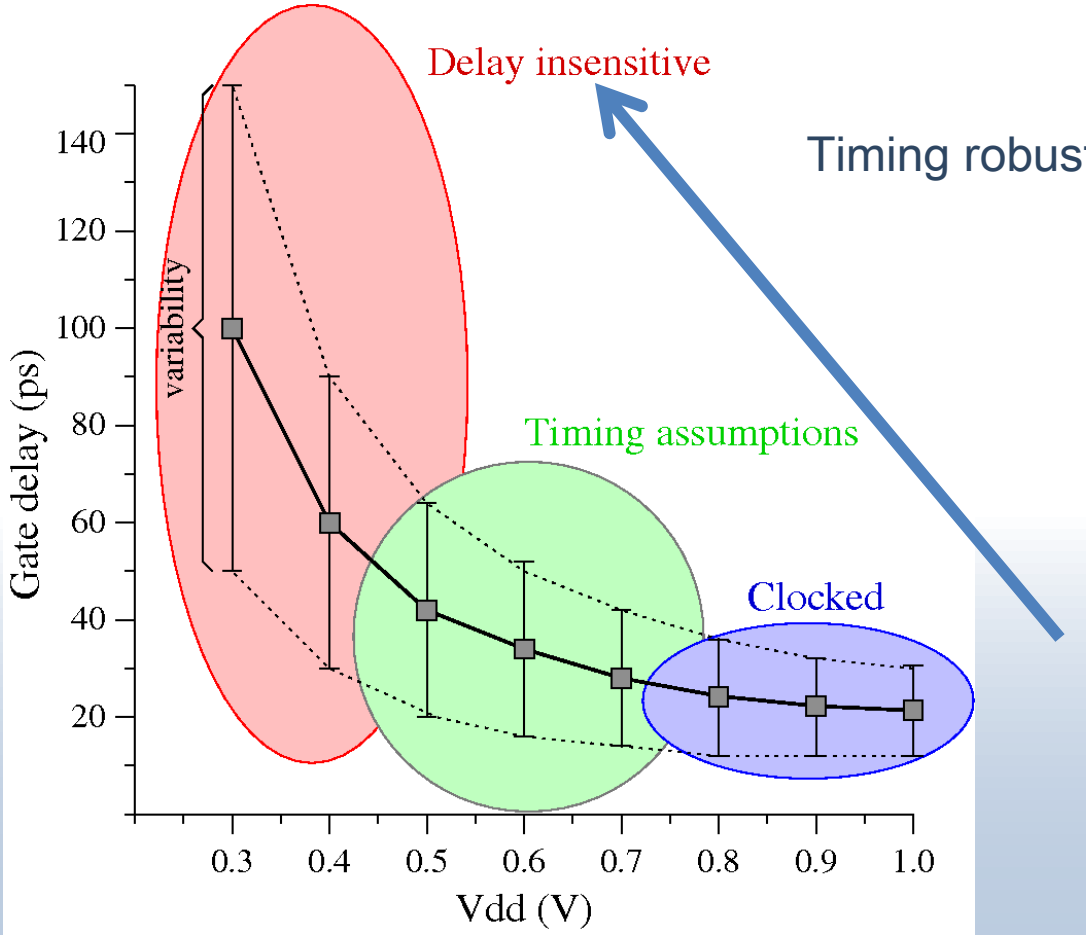

Technology node:
90nm

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

Towards designing 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?

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

PLUS Cross-level adaptation to uncertainty and errors

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)

Asynchronous (elastic, self-timed) design principles improve robustness to uncertainty and efficiency of both sensing and control in adaptation process

- Process-level adaptation to uncertainty and errors

Cell-level adaptation example: robust synchronizer

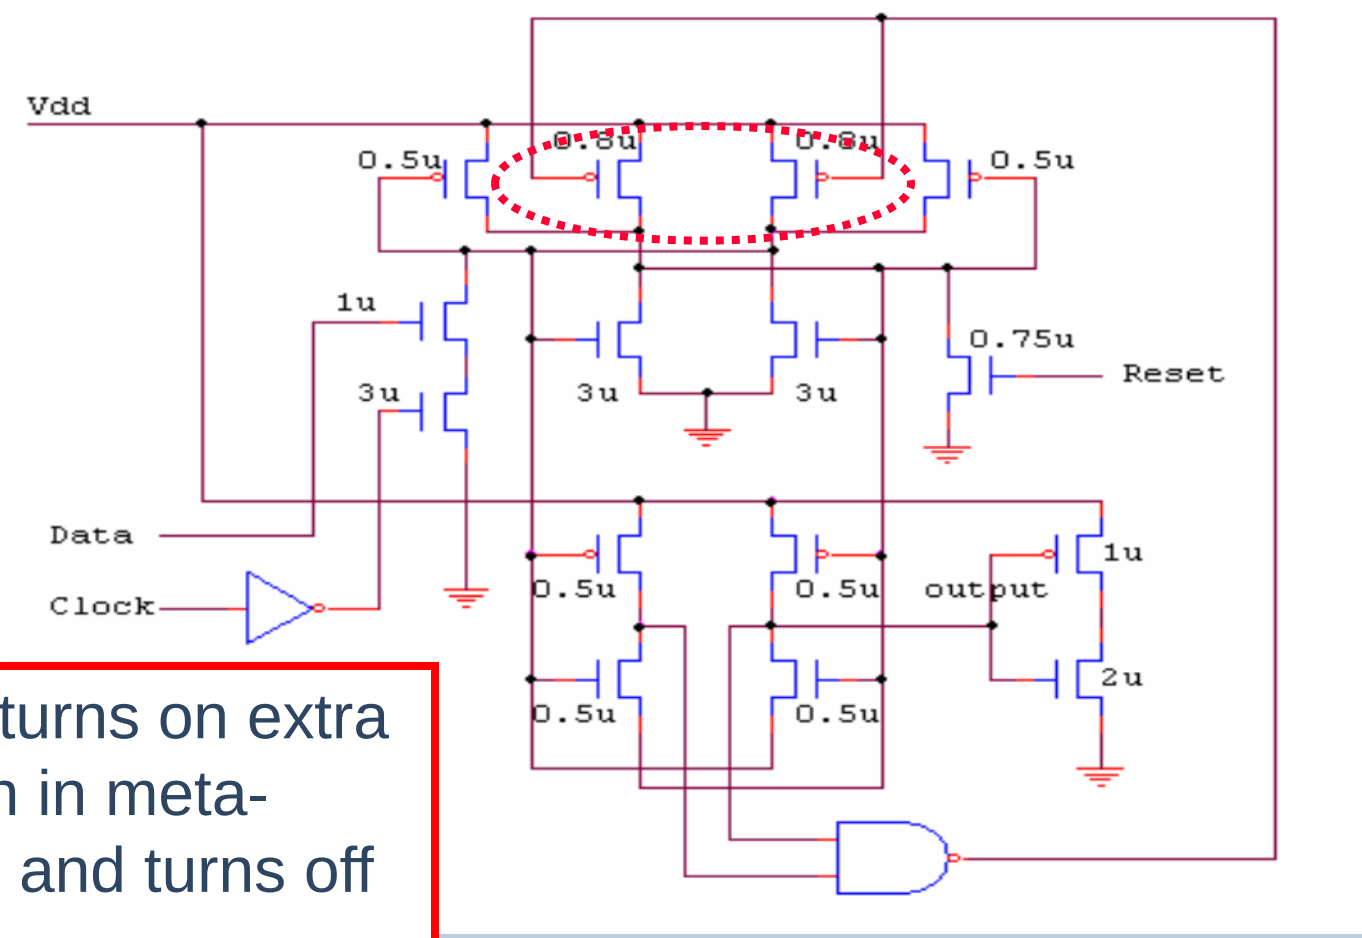

This circuit turns on extra power when in meta-stable state and turns off after that

Circuit-level: speed-independent SRAM

- **Mismatch between delay lines and SRAM memories when reducing Vdd**

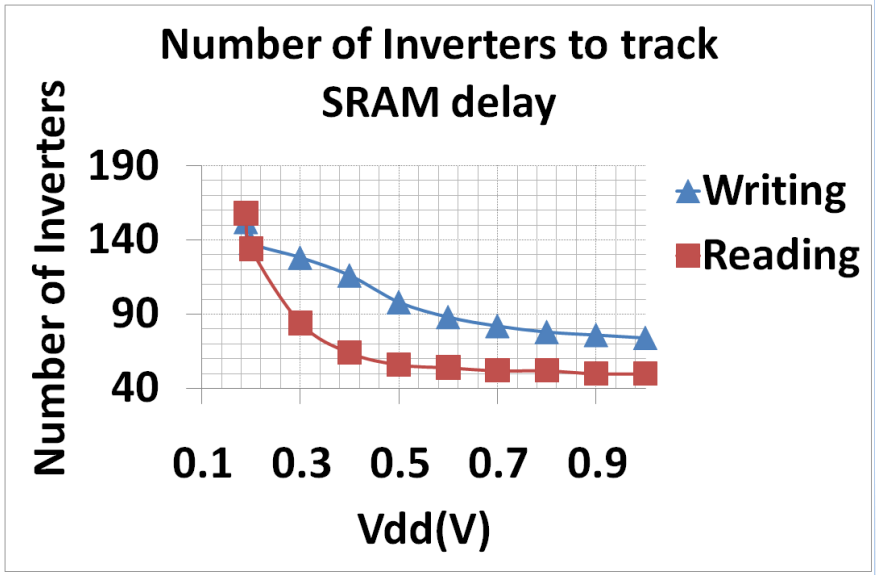

For example, under 1V Vdd, the delay of SRAM reading is equal to 50 inverters and under 190mV, the delay is equal to 158 inverters

- **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?!**

A. Baz et.al. PATMOS 2010, JOLPE 2011

Conclusions

- **Future systems will increasingly be energy-modulated**
- **Energy modulated systems are power-proportional and power-efficient, they are also power-adaptive**
- **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 Vdd 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**
  - **Cross-layer management of uncertainty and errors in computation**

Acknowledgements

Members of “Microelectronics Systems Design” research group at Newcastle:
Andrey Mokhov, Delong Shang, Danil Sokolov, Fei Xia, Reza Ramezani, Abdullah Baz, Xuefu Zhang, involved in the Power-Adaptive Computing research

Members of the Holistic Project consortium

EPSRC funding:

- Holistic Project “Next Generation Energy-Harvesting Electronics: A Holistic Approach”, involving Universities of Southampton, Bristol, Newcastle and Imperial College
- Dream Fellowship “Energy-Modulated Computing” (2011-2013)

Example 2: self-powered console

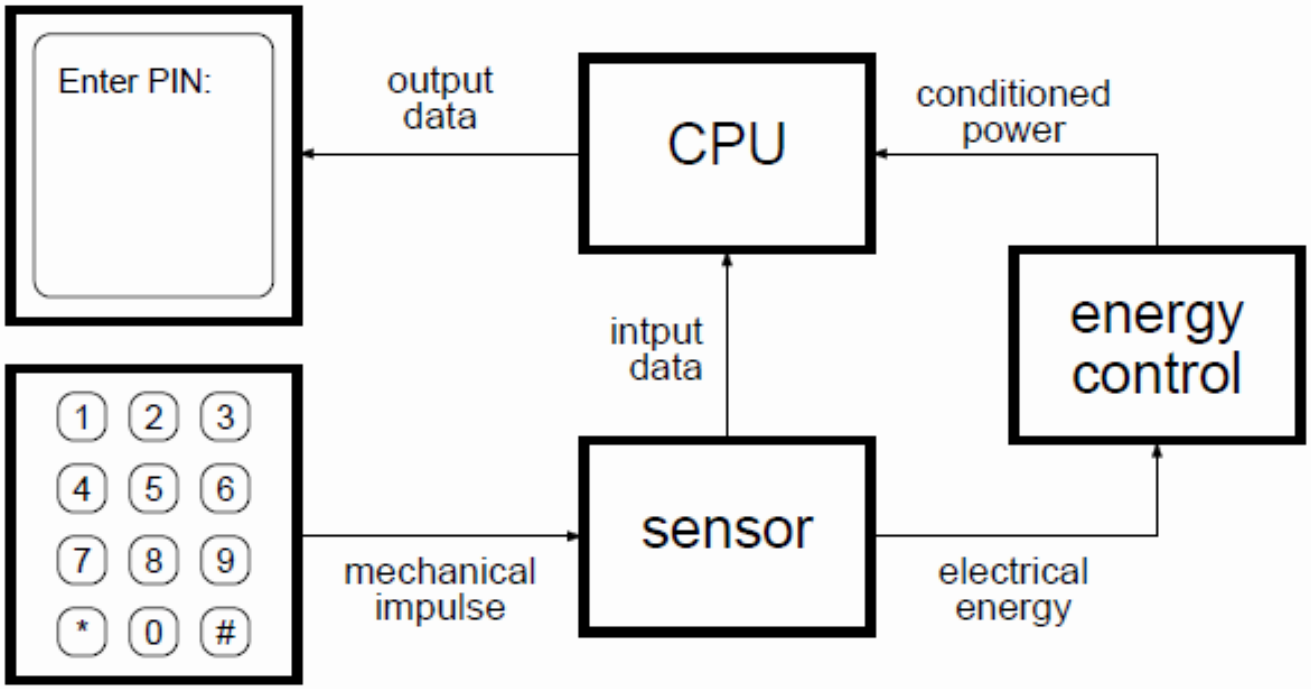

- Polling/survey data acquisition devices
- PIN-protected electronic lock
- Word processor

Energy-efficiency vs Robustness