An Architecture for a Power-Aware Distributed Microsensor Node

Rex Min, Manish Bhardwaj, Seong-Hwan Cho, Amit Sinha, Eugene Shih, Alice Wang, and Anantha Chandrakasan
Massachusetts Institute of Technology

Application Driver

Distributed Microsensor Networks

- *Ad hoc networking*
- *Collaborative data gathering*
- *Robustness & fault-tolerance*

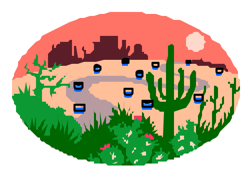

Uniqueness:

- Event driven: **Low duty cycles**
- Low bandwidth: **bits/sec to kbits/sec**
- High Spatial Density: **0.1-20 nodes/m²**
- Short Tx distances: **5-10m typ. (< 100m)**
- Remote deployment: can't swap battery

Diversity:

Node Roles:	*Environment:*	*User Demands:*
Sensor	Event arrival rate/type	Tolerable latency
Relay	Ambient noise	Result SNR
Data aggregator	Signal statistics	Pr(Detection)

How to achieve months to years of life from a single battery?

Platform Architecture

A Node that Illustrates Power-Aware Methodologies

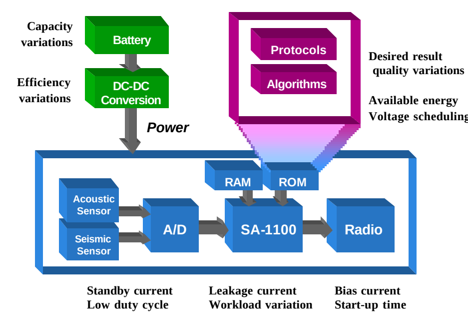

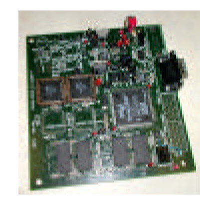

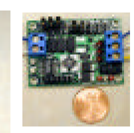

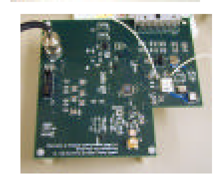

Above: miniaturized DVS circuit

Upper-left: sensor and processor board.

Left: radio baseband prototype.

Every component of the system considers power-awareness. First generation prototypes are with COTS components.

POWER-AWARE METHODOLOGIES

"Low-Power" differs from "Power-Aware"

Low-Power

- Component-by-component optimization
- Reduction of worst-case power dissipation

Power-Aware

- Graceful energy scalability across a diversity of operating conditions
- Energy-quality trade-offs
- Collaboration across levels of the system hierarchy

Application drives methodology:
High diversity => graceful scalability

Formalize Operational Diversity

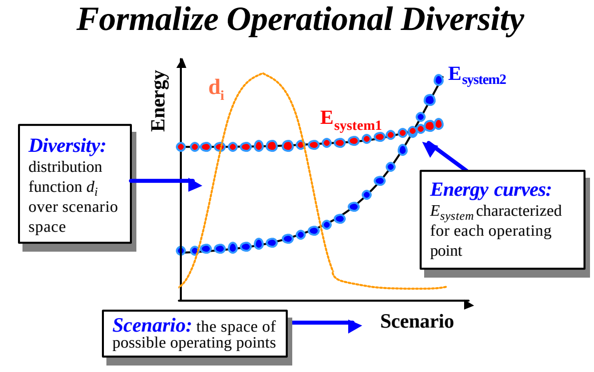

Idle Mode Issues

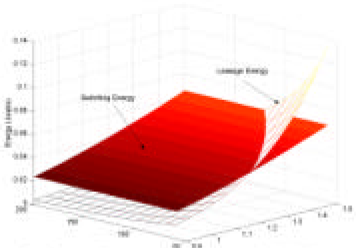

Processor: Leakage vs. Switching Energy

Low-power design has traditionally focused on switching energy, given by

$$E_{switch} = C_{tot} V_{DD}^2$$

Measurements from the SA-1100 indicate that the leakage energy, given by

$$E_{leak} = V_{DD} (I_0 e^{\frac{V_{DD}}{nV_T}}) t$$

can actually exceed switching energy in certain modes of operation. As shown below, this leakage component can not be ignored for a microsensor node.

Leakage dominates for low duty cycles.

The SA-1100 conserves power by entering a low-power *idle* mode. Idling the processor reduces switched capacitance, but a substantial amount of leakage current continues to flow. At duty cycles below 20%, leakage energy exceeds switching energy. The low duty cycle characteristics inherent to a sensor network demand an awareness of leakage through OS-directed selective component shutdown.

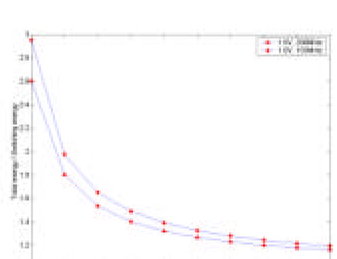

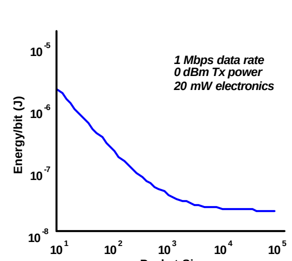

Radio: Startup energy exceeds transmit energy.

The low data rate and short transmission distance inherent to microsensor networks implies short, intermittent packets. In this regime, the energy required to start up a radio far exceeds the energy of the actual transmission. Until ultra-low startup time radios are devised, communications should be buffered to the point of maximum tolerable latency.

Active Mode Issues

Circuit-level: Dynamic Voltage Scaling (DVS)

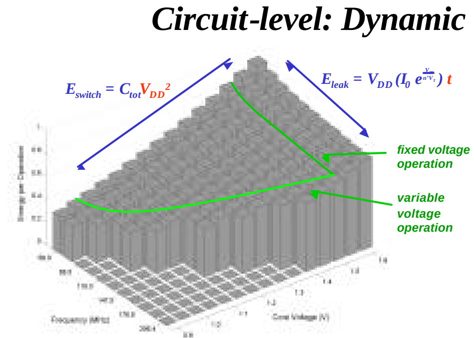

Change voltage with frequency.

Reducing the clock frequency enables the static core of the SA-1100 to run at a lower voltage. These measurements show how voltage scaling reduces the energy consumed *per operation* on the SA-1100.

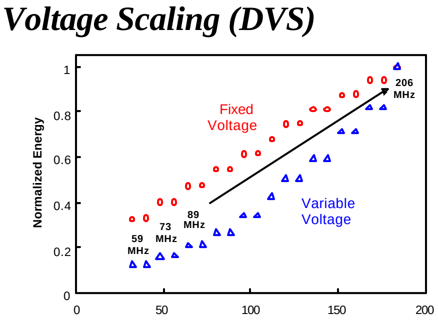

Energy-scalable digital filter with DVS.

A variable-tap FIR filtering algorithm imposes operational diversity on the SA-1100. The µOS scales the clock frequency for maximum processor utilization. When voltage is scaled as well, energy savings can exceed 50%.

Algorithm-level: Energy-Quality Scalability

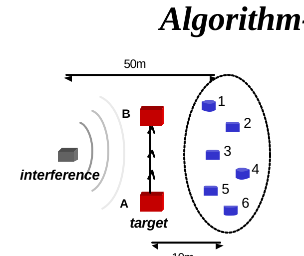

Beamforming example.

Beamforming combines data from several sensors into a single, high-SNR signal. In this scenario, a line of sensors tracks a target moving from A to B. Background interference limits SNR.

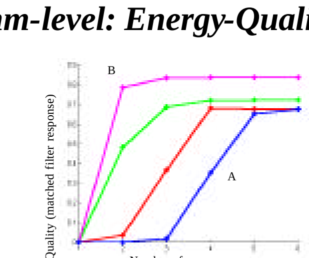

Poor E-Q scalability.

How to choose *M* sensors for beamforming? Any fixed order (#1-#6 in the example above) causes the energy-quality curves to vary substantially as the target moves.

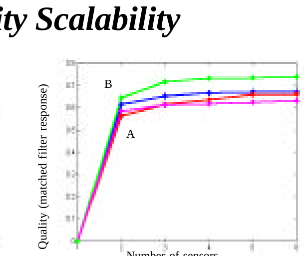

Improving scalability.

A *most significant first* transform ranks the signals via a frequency domain SNR estimation and quicksort. With a worst-case overhead of 0.44% on the SA-1100, the curves now refine more consistently.