Approximating **Outside the Processor**

Phillip **Stanley-Marbell** and Martin **Rinard**
psm@mit.edu

Massachusetts Institute of Technology

WAX'15, co-located with PLDI'15, 13th June 2015

Approximating **Outside the Processor**

(Source: Fitbit)

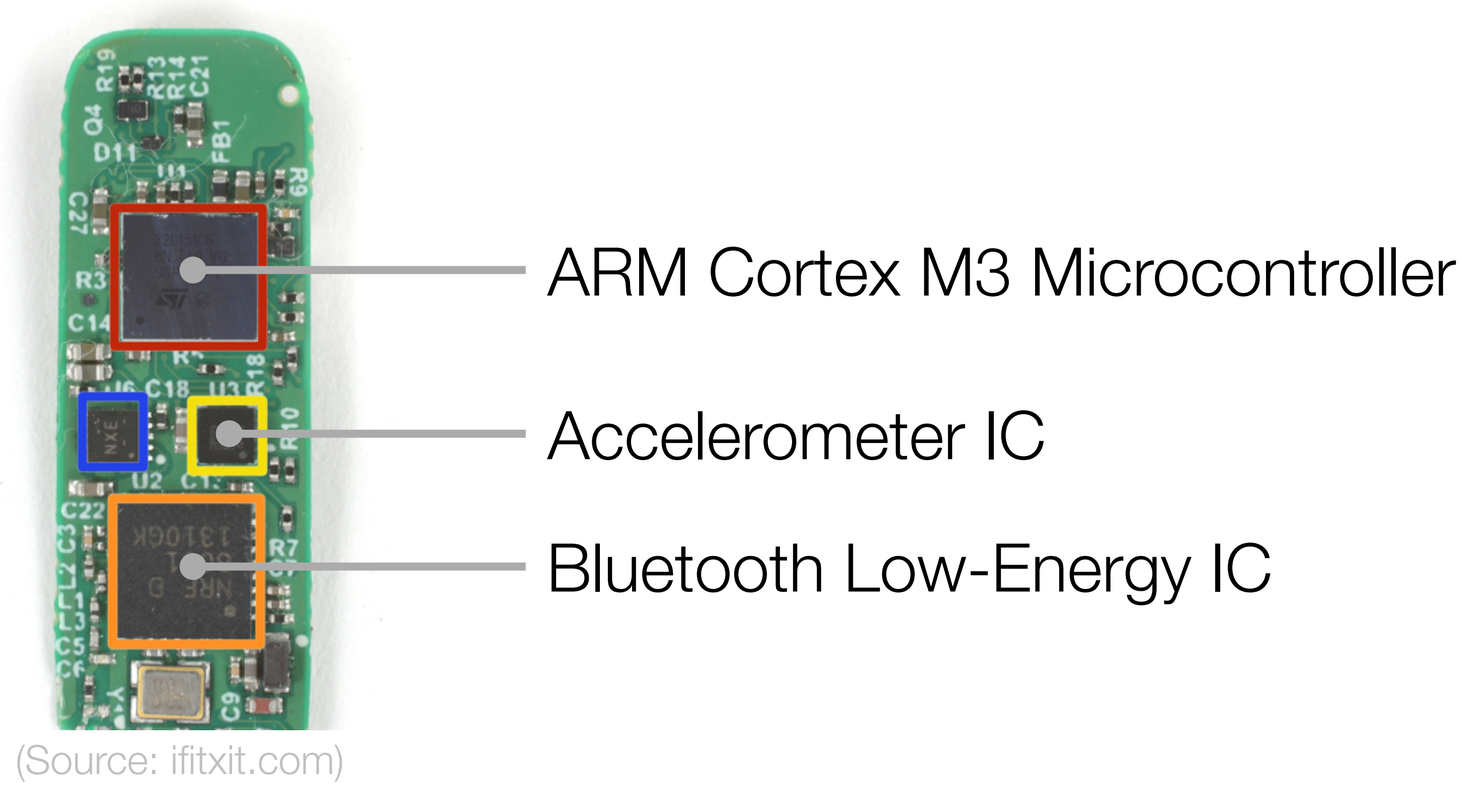

(Source: ifitxit.com)

Approximating **Outside the Processor**

(Source: Samsung)

320x320 AMOLED Display, 24 bits per pixel color

Dual-Core ARM (Exynos 3250)

ARM Cortex M4 Microcontroller

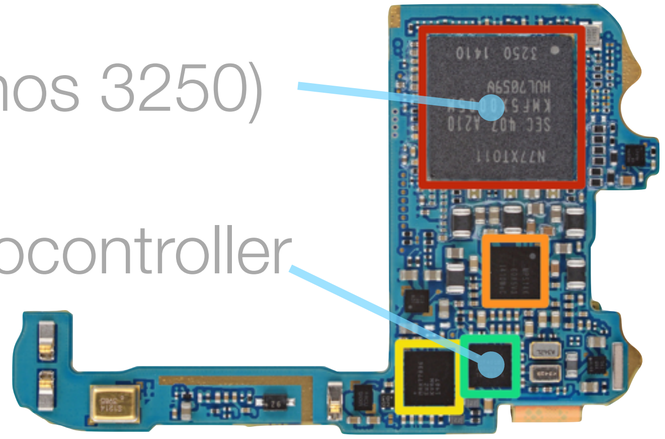

(Source: ifitxit.com)

Sensors/Displays Dominate Power Dissipation in Wearables

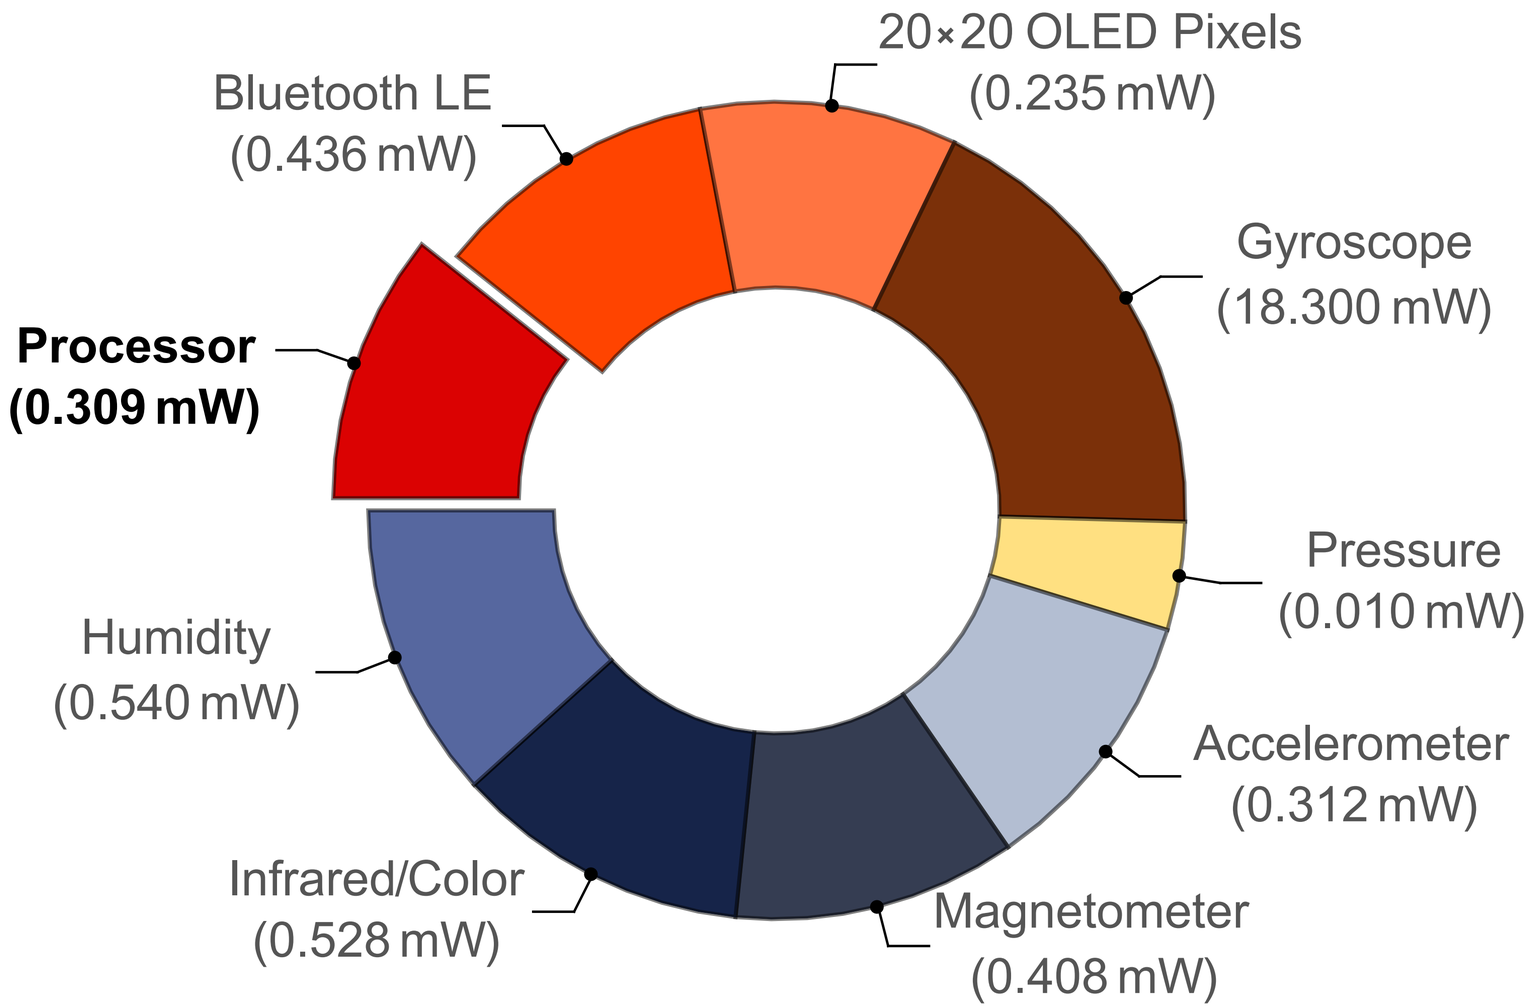

(Note: sectors scaled logarithmically due to large range of values)

To improve system-wide energy-efficiency, **focus on the dominant fraction**

Approximating Outside the Processor

We make three arguments, in light of these observations

For sensor-driven system such as wearables and “internet-of-(every)things”

❶ Explore approximation in sensors / inputs

❷ Explore approximation in displays / outputs / communication

Across all system types:

❸ What are bounds on benefit from reducing precision, accuracy, reliability/certainty?

(If best-case achievable gains are small, why bother…)

❶ Approximate Inputs

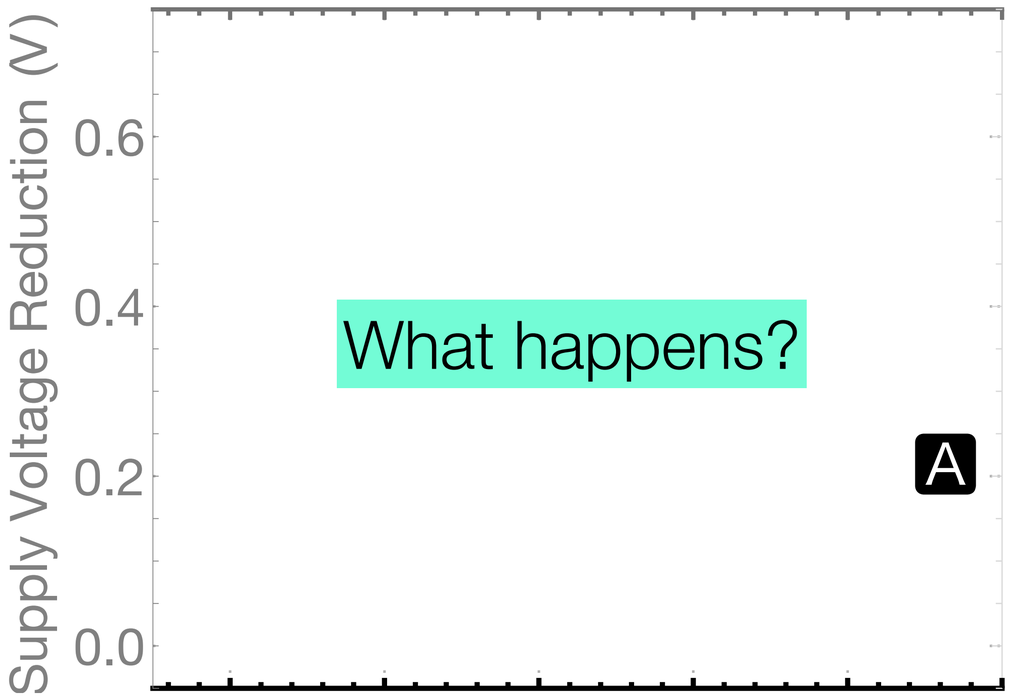

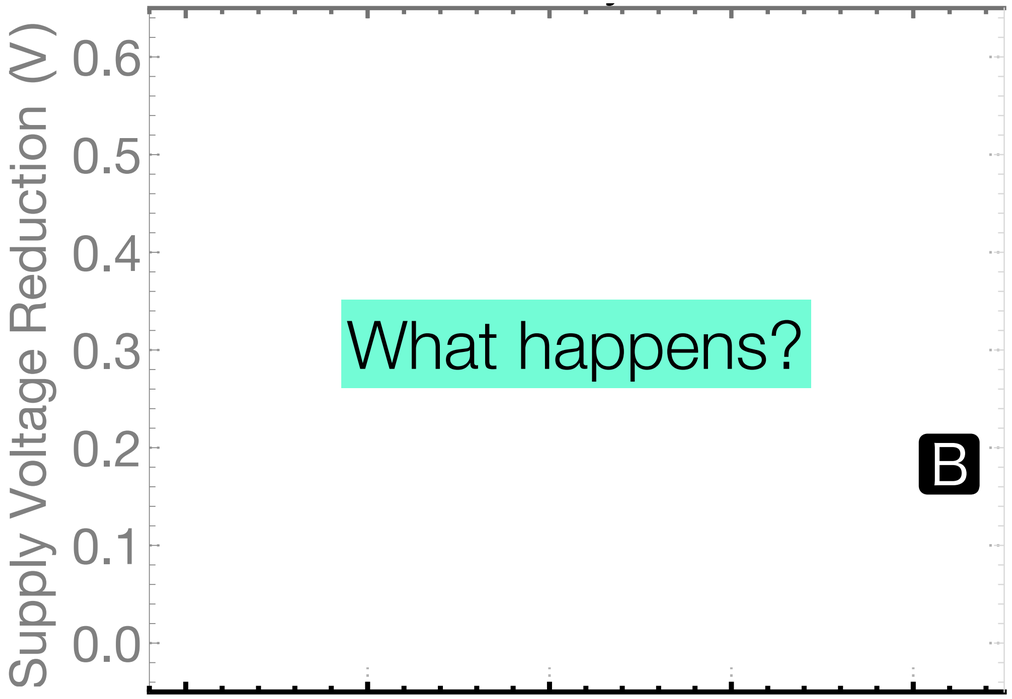

A

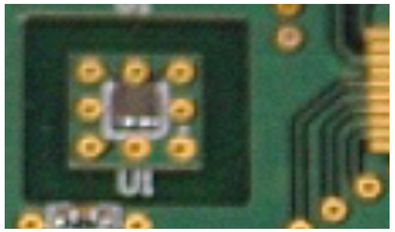

TI TMP006
IR Temperature
Sensor

B

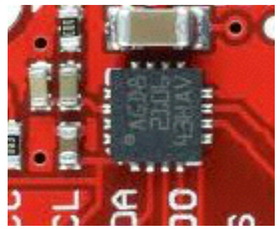

ST L3G4200D
Gyroscope
Sensor

P. Stanley-Marbell and M. Rinard. “Lax: Driver Interfaces for Approximate Sensor Device Access”, *USENIX HotOS’15*, 2015

❶ Approximate Inputs

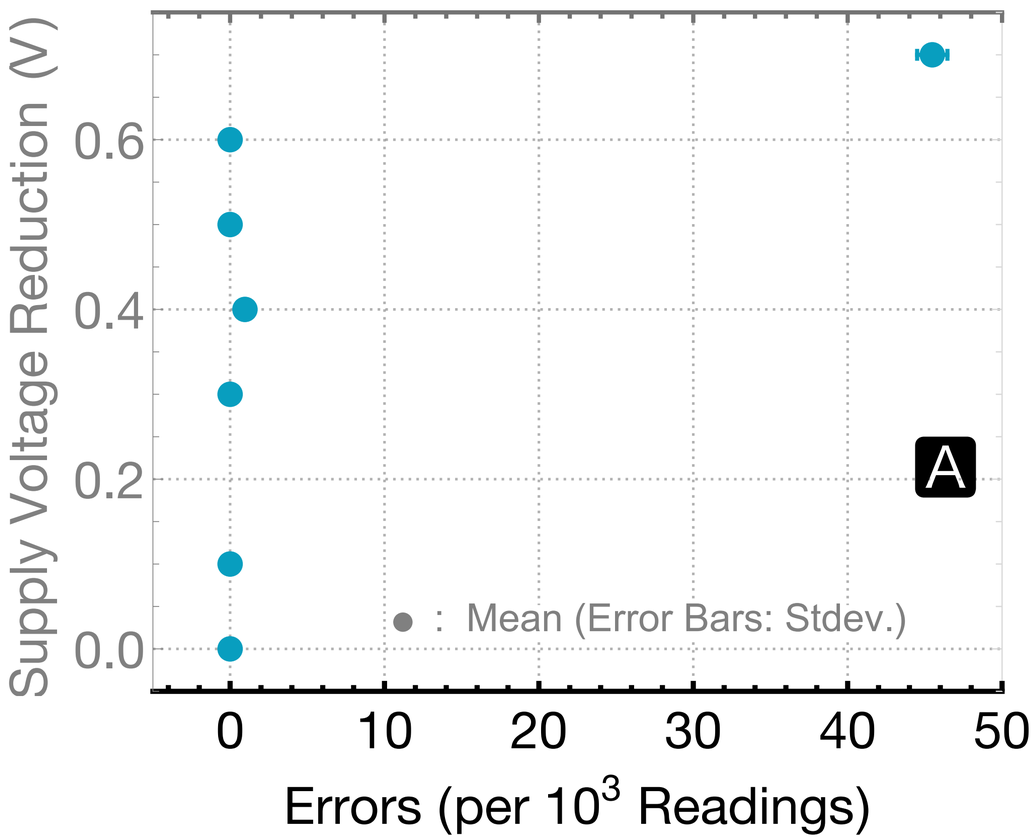

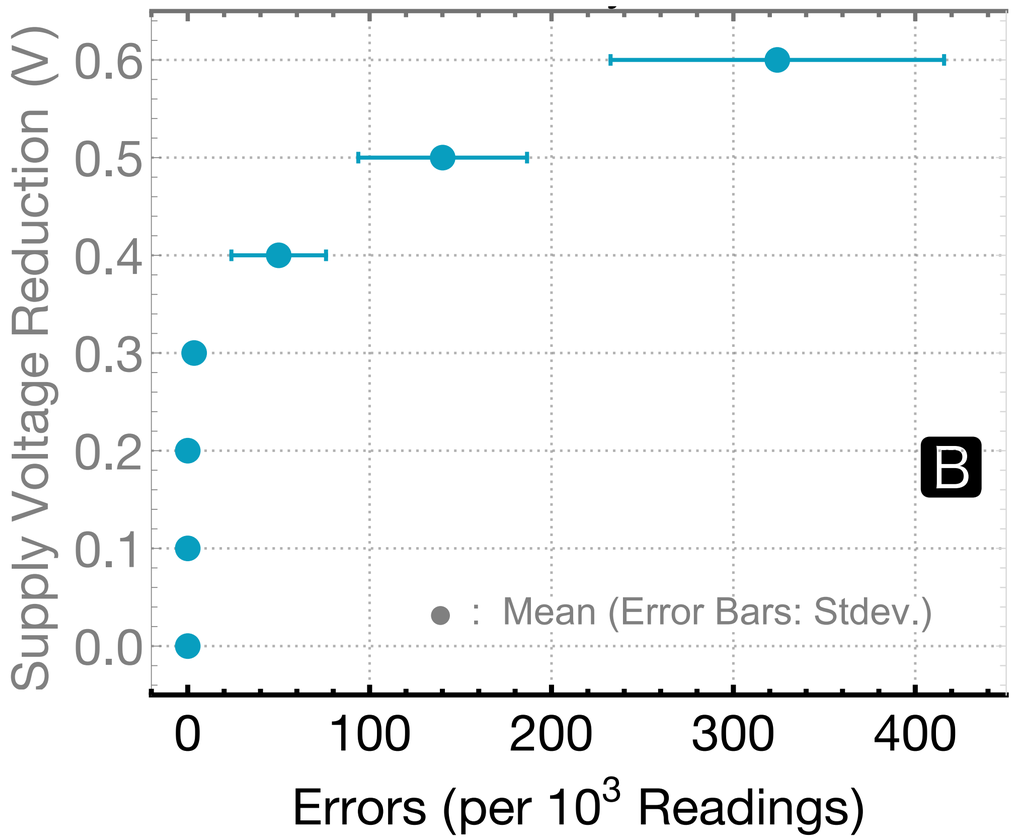

A

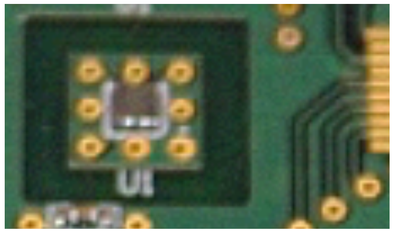

TI TMP006
IR Temperature
Sensor

B

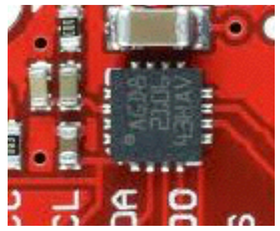

ST L3G4200D
Gyroscope
Sensor

P. Stanley-Marbell and M. Rinard. “Lax: Driver Interfaces for Approximate Sensor Device Access”, *USENIX HotOS’15*, 2015

❶ Approximate Inputs

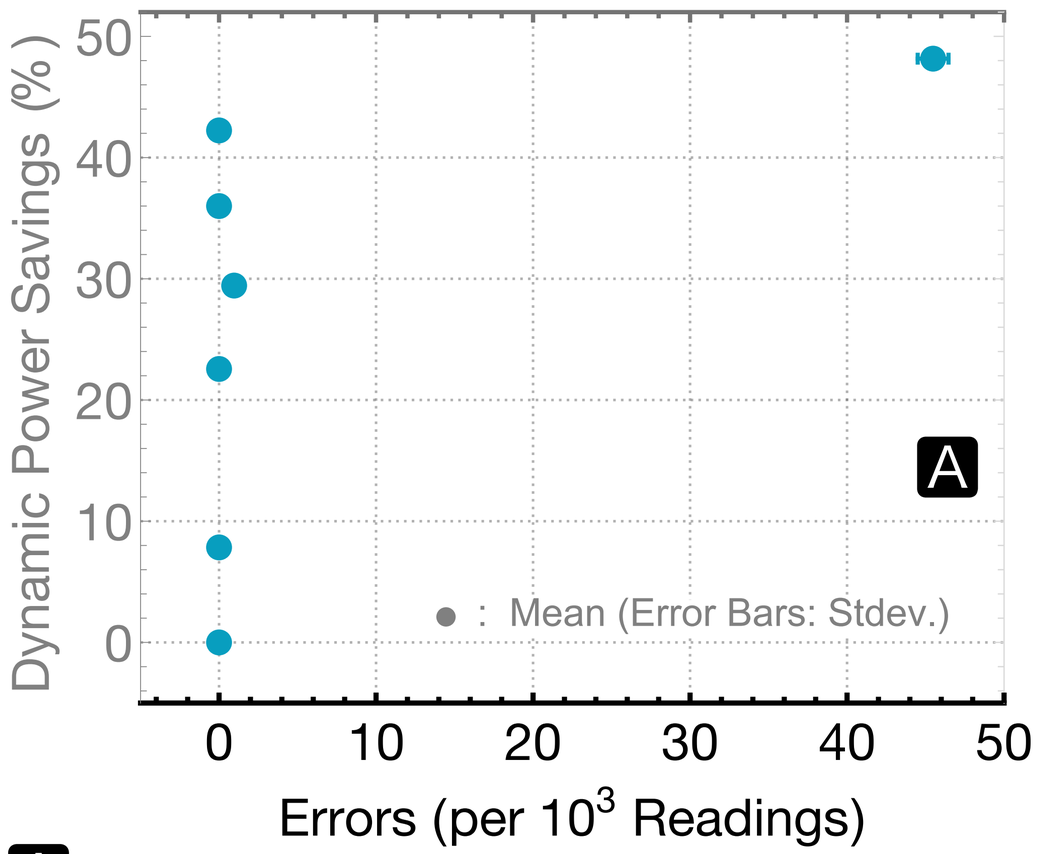

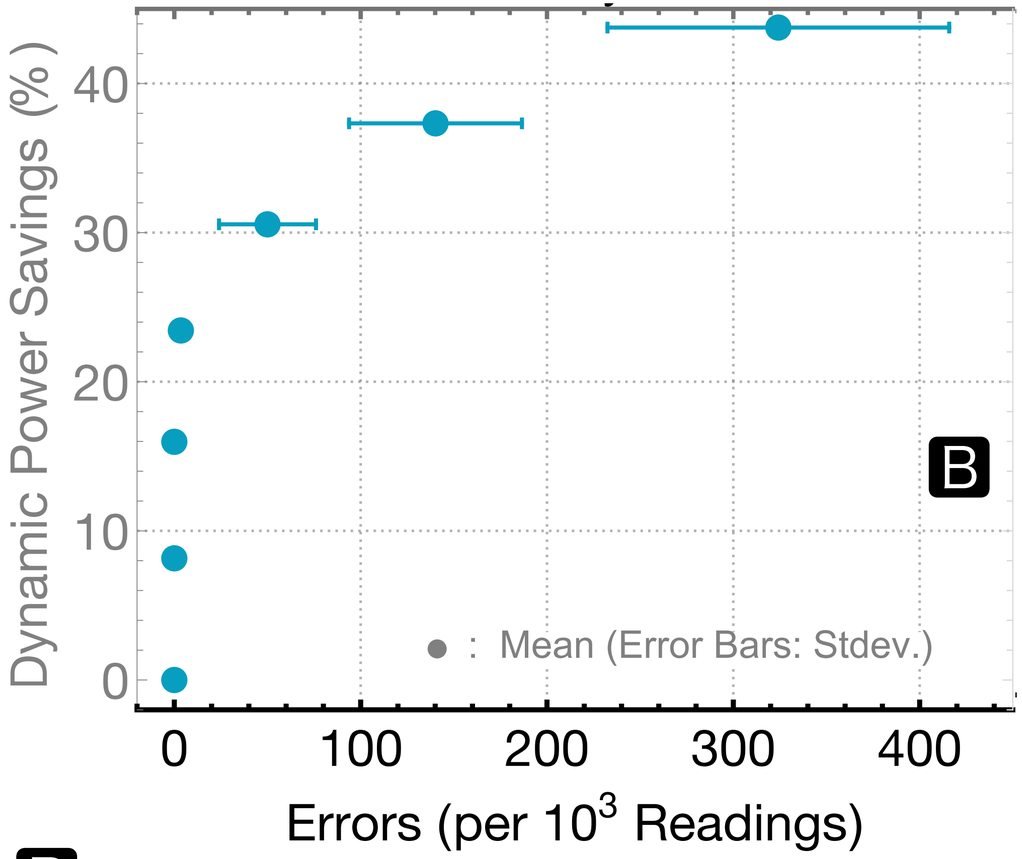

A

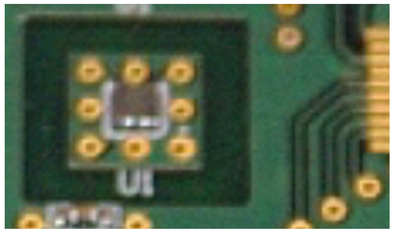

TI TMP006
IR Temperature
Sensor

B

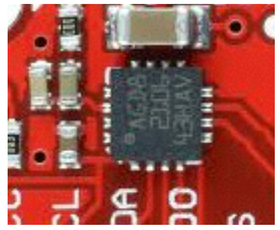

ST L3G4200D
Gyroscope
Sensor

P. Stanley-Marbell and M. Rinard. “Lax: Driver Interfaces for Approximate Sensor Device Access”, *USENIX HotOS’15*, 2015

❷ Approximate Outputs / Communication

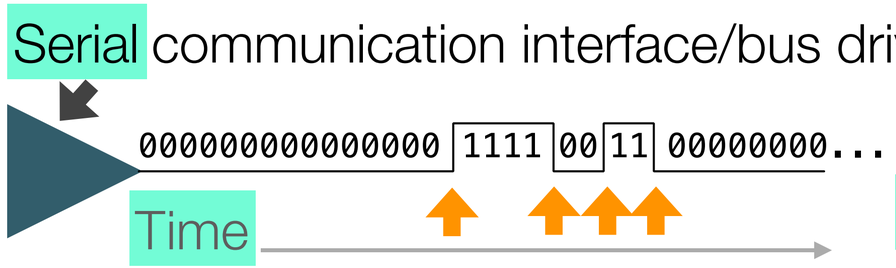

More transitions : more dynamic power dissipation

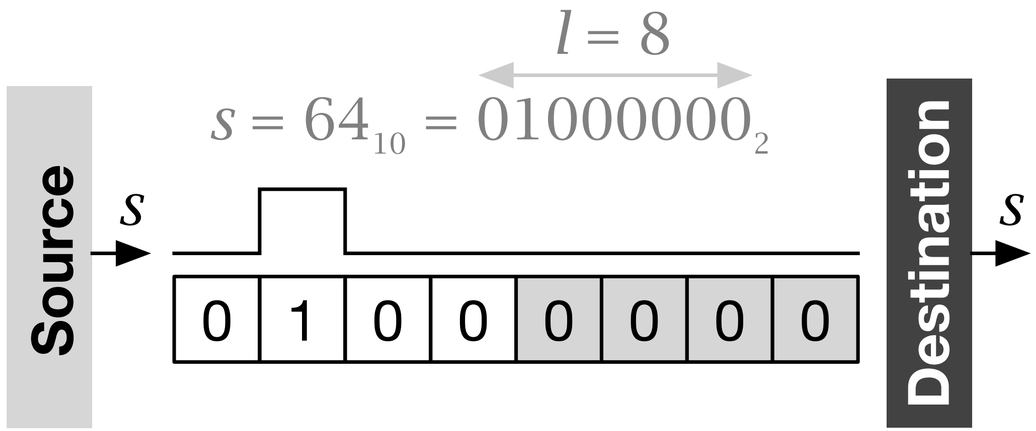

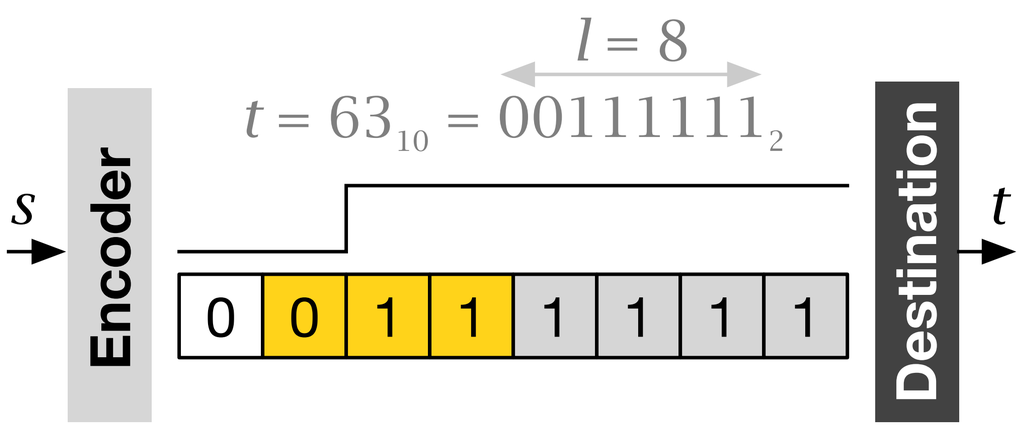

P. Stanley-Marbell and M. Rinard. “Value-Deviation-Bounded Serial Data Encoding for Energy-Efficient Approximate Communication”, *MIT-CSAIL-TR-2015-022*, 2015

❸ Bounds: How Badly (and Often) Do Values Deviate?

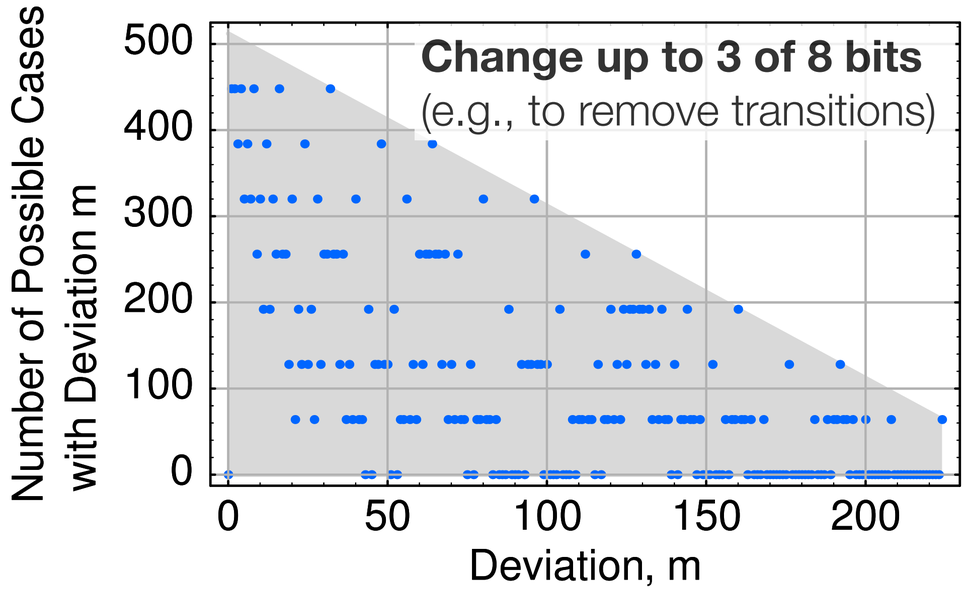

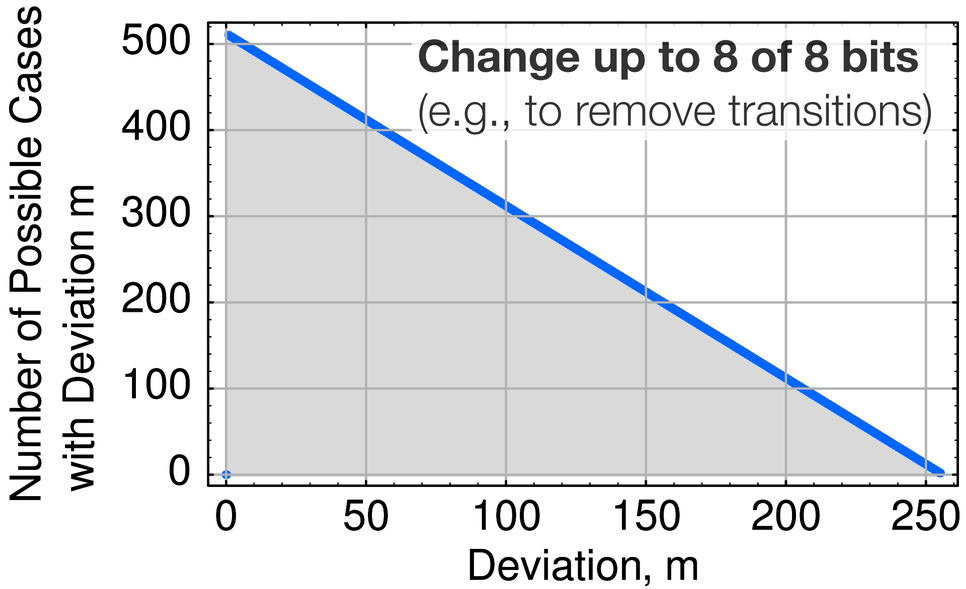

Count of cases is given by number of solutions to Diophantine equation pair:

(*w* and *v* are two *L*-bit words)

$$\left| \sum_{i=0}^{L-1} w_i 2^i - \sum_{i=0}^{L-1} v_i 2^i \right| = m$$

Deviation from correctness

$$\sum_{i=0}^{L-1} (w_i(1 - v_i) + v_i(1 - w_i)) = k$$

Number of bits we perturb

An **upper bound on number of cases is** $2^{L+1} - 2m$ (shaded gray region in plots above)

P. Stanley-Marbell “Encoding Efficiency of Digital Number Representations under Deviation Constraints”, *IEEE Information Theory Workshop, ITW’09, 2009*

Position Statement, Conclusions, Q&A

Approximating Outside the Processor

❶ We should **focus on the domains that will be important** in the future

❷ For the important domains, we should **focus on bottlenecks**

❸ For the growing domain of IoT/wearables: **sensors, displays, communication**

❹ **Understanding upper limits of benefits** informs choice of realistic techniques

Position Statement, Conclusions, Q&A

Approximating Outside the Processor

1. We should **focus on the domains that will be important** in the future
2. For the important domains, we should **focus on bottlenecks**
3. For the growing domain of IoT/wearables: **sensors, displays, communication**
4. **Understanding upper limits of benefits** informs choice of realistic techniques