

Hardware-Software Inexactness in Noise-aware Design of Low-Power Body Sensor Nodes

*Hossein Mamaghanian**

Giovanni Ansaloni⁺

Mohamed M. Sabry⁺

David Atienza⁺

*Pierre Vandergheynst**

**Signal Processing Laboratory (LTS2)*

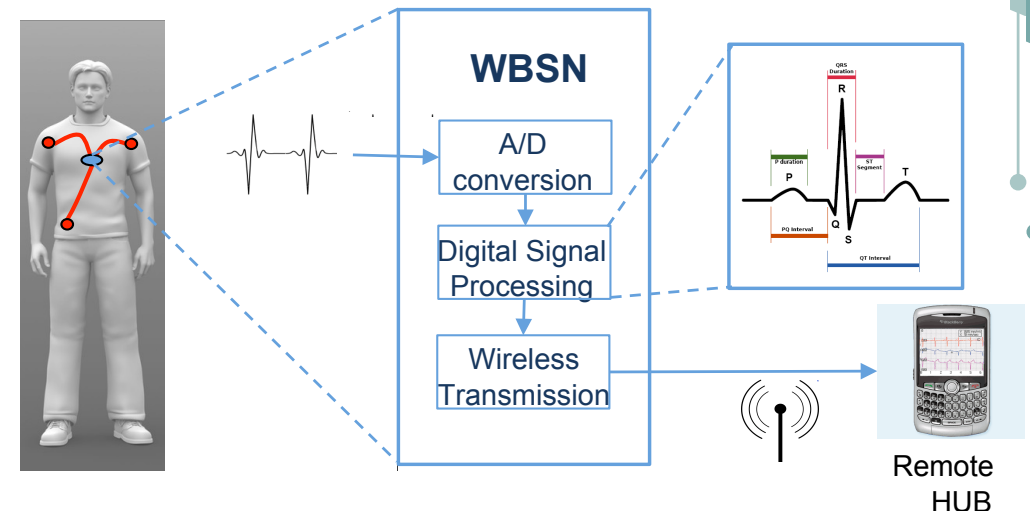
+Embedded Systems Lab (ESL)

*École Polytechnique Fédérale de
Lausanne (EPFL), Switzerland*

Designing with Uncertainty - Opportunities & Challenges workshop
York, UK, 17-19 March 2014

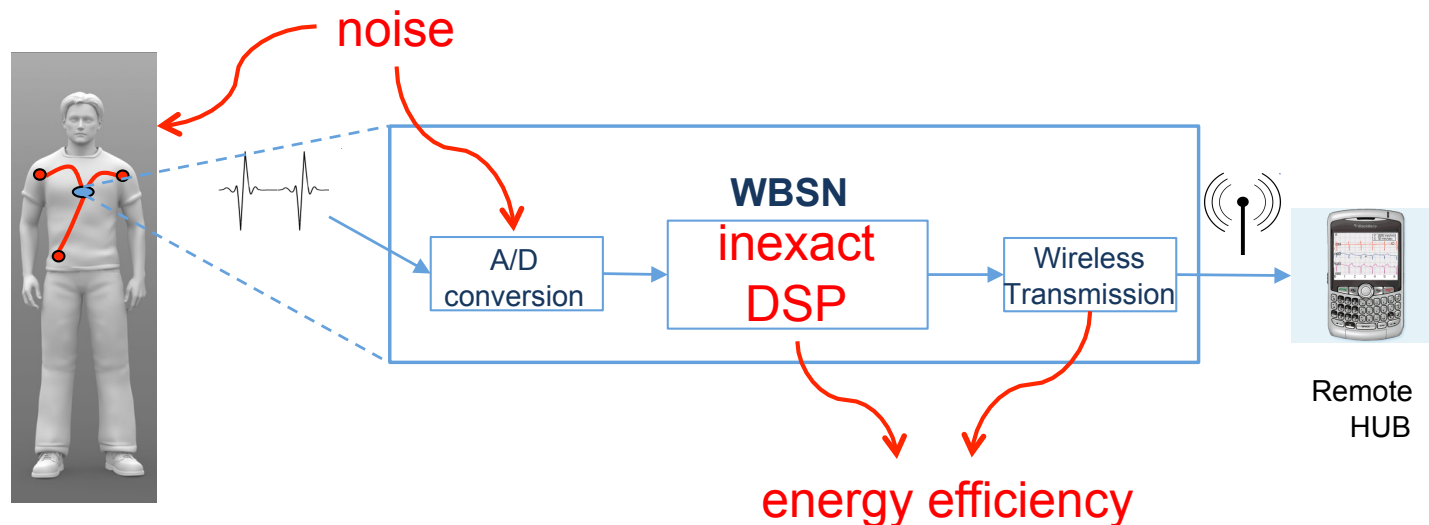
Wireless Body Sensor Nodes (WBSNs)

- Continuous monitoring of bio-signals
 - Respiration
 - Blood Flow
 - ECG
- Three phases
 - acquisition
 - processing
 - wireless transmission
- **Energy efficiency is key**
 - wearability
→ small battery
 - continuous monitoring



Noise in WBSN signal acquisitions

- Noise is un-avoidable in WBSN systems
 - external noise (for ECG: muscle contractions, respiration...)
 - A/D conversion noise
 - Relevant data is inherently noise resilient
 - qualitative and/or statistical
- opportunity for inexact, energy efficient DSP



Inexact DSP in WBSNs

- **Software:** on-node lossy compression
 - small impact on the quality-of-service
 - reduction in wireless transmission bandwidth
 - energy efficiency in wireless link
 - Example: Compressed Sensing

- **Hardware:** significance-based computing
 - allows errors in computations contributing the least to results
 - only most-significant computation is error-free
 - energy reduction for embedded DSP
 - Example: DWT compression

Inexact DSP in WBSNs

- **Software:** on-node lossy compression
 - small impact on the quality-of-service
 - reduction in wireless transmission bandwidth
 - energy efficiency in wireless link
 - Example: Compressed Sensing

- **Hardware:** significance-based computing
 - allows errors in computations contributing the least to results
 - only most-significant computation is error-free
 - energy reduction for embedded DSP
 - Example: DWT compression

Compressed Sensing at a glance

- Bio-signals are highly sparse in a transformed domain
 - they can be represented with few (approximately) non-zero coefficients
 - ECG is sparse in the wavelet domain
- **CS-based compression**: collection of $M (\ll N)$ linear random measurements

$$\begin{array}{c}
 \text{Measurement vector} \rightarrow \mathbf{y}_{M \times 1} = \mathbf{\Phi}_{M \times N} \cdot \mathbf{x}_{N \times 1} \\
 \begin{array}{l}
 \text{Measurement/Sensing matrix} \\
 \text{(Gaussian random matrix)} \\
 \text{Original ECG vector}
 \end{array}
 \end{array}$$

- **CS reconstruction**: convex optimization problem (at the receiver side!)

$$\min_{\tilde{\alpha} \in \mathcal{R}^N} \|\tilde{\alpha}\|_1 \quad \text{Subject to:} \quad \|\Phi\Psi\tilde{\alpha} - y\|_2 \leq \sigma$$

CS of noisy signals

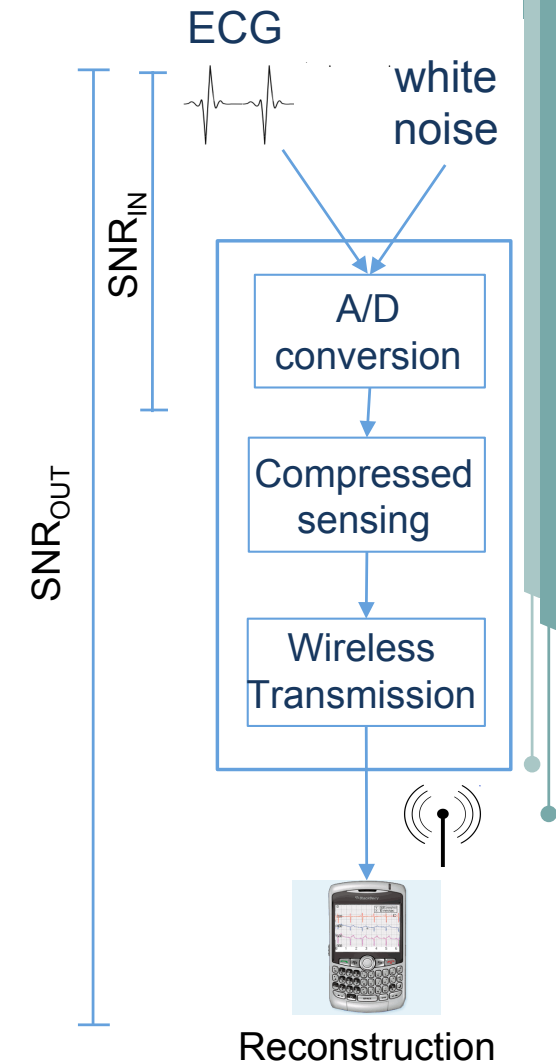

- Compression ratio ($CR = M/N$) can be adapted to the quality of acquired signals
 - external and A/D noiseand/or the desired quality of the output
 - “good” level for ECGs: SNR of 21db [1]
- Experimental database:
 - MIT-Arrhythmia ECG Database [2] with different levels of white gaussian noise

[1] Y. Zigel, A. Cohen, and A. Katz, “The weighted diagnostic distortion (wdd) measure for ecg signal compression,” Biomedical Engineering, IEEE Transactions on, vol. 47, no. 11, pp. 1422–1430, Nov. 2000.

[2] www.physionet.org/physiobank/


Experimental results

- $SNR_{IN} = 20 \times \log_{10}(\|x_{IN}\|_2 / \|n\|_2)$
- $SNR_{OUT} = 20 \times \log_{10}(\|x_{out}\|_2 / \|x_{IN} - x_{out}\|_2)$




Experimental results

- low $\text{SNR}_{\text{IN}} \rightarrow \text{SNR}_{\text{OUT}}$ limited by input noise
 - increasing CR doesn't affect output quality (but requires energy for DSP and transmission!)



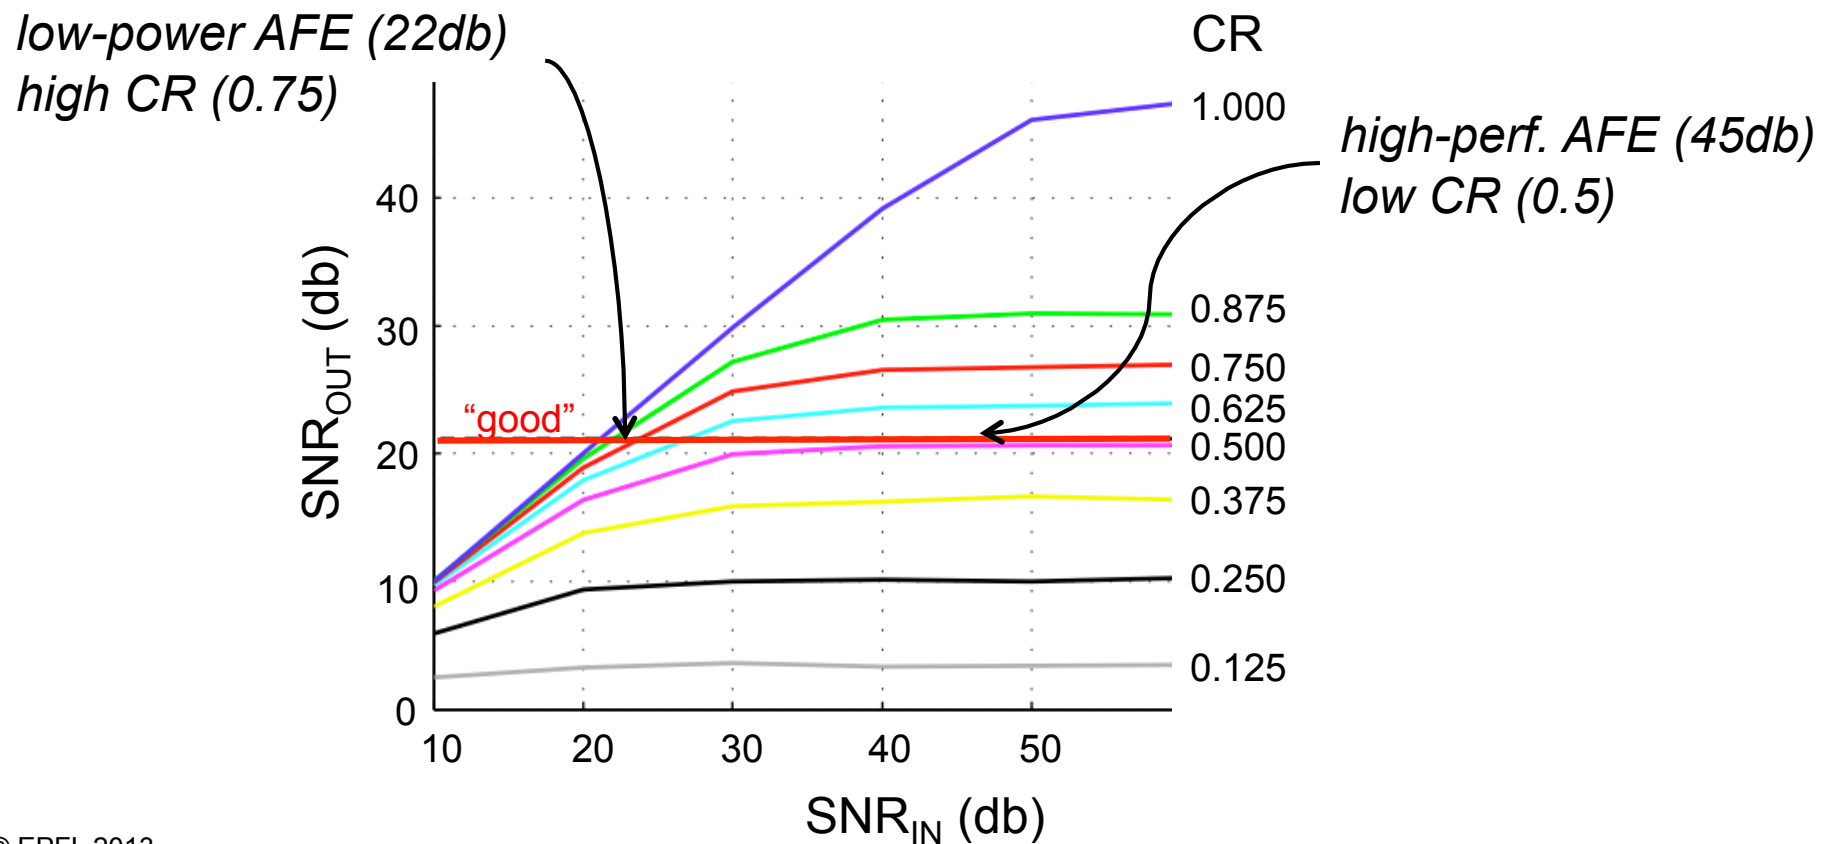
Experimental results

- high $\text{SNR}_{\text{IN}} \rightarrow \text{SNR}_{\text{OUT}}$ limited by CR
 - increasing SNR_{IN} doesn't affect output quality (but requires resources for filtering/sampling!)



Experimental results

- trade-off between sampling, DSP and transmission for a fixed SNR_{OUT}

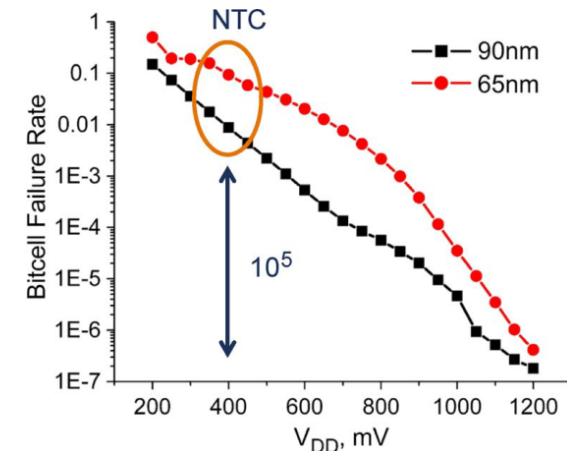


Inexact DSP in WBSNs

- Software: on-node lossy compression
 - small impact on the quality-of-service
 - reduction in wireless transmission bandwidth
 - energy efficiency in wireless link
 - Example: Compressed Sensing
- **Hardware**: significance-based computing
 - allows errors in computations contributing the least to results
 - only most-significant computation is error-free
 - energy reduction for embedded DSP
 - Example: DWT compression

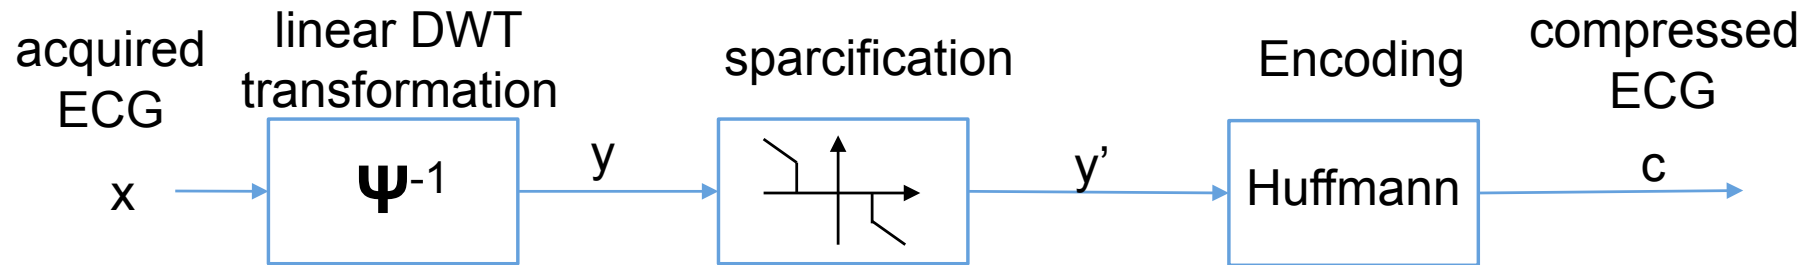
Significance-based computing at a glance

- Near threshold (low-V_{DD}) computing
 - Enhances energy efficiency
 - Compromises data integrity in memories
- Protecting all data is costly
 - additional energy required for Hw and Sw implementations
- In most WBSN applications, computations do not contribute equivalently to the QoS
- **Significance-based computing**
 - Critical computations: guaranteed protection
 - Non-critical: best-effort protection




Dreslinski et al. IEEE, 2010

Case study: Embedded DWT compression of ECGs

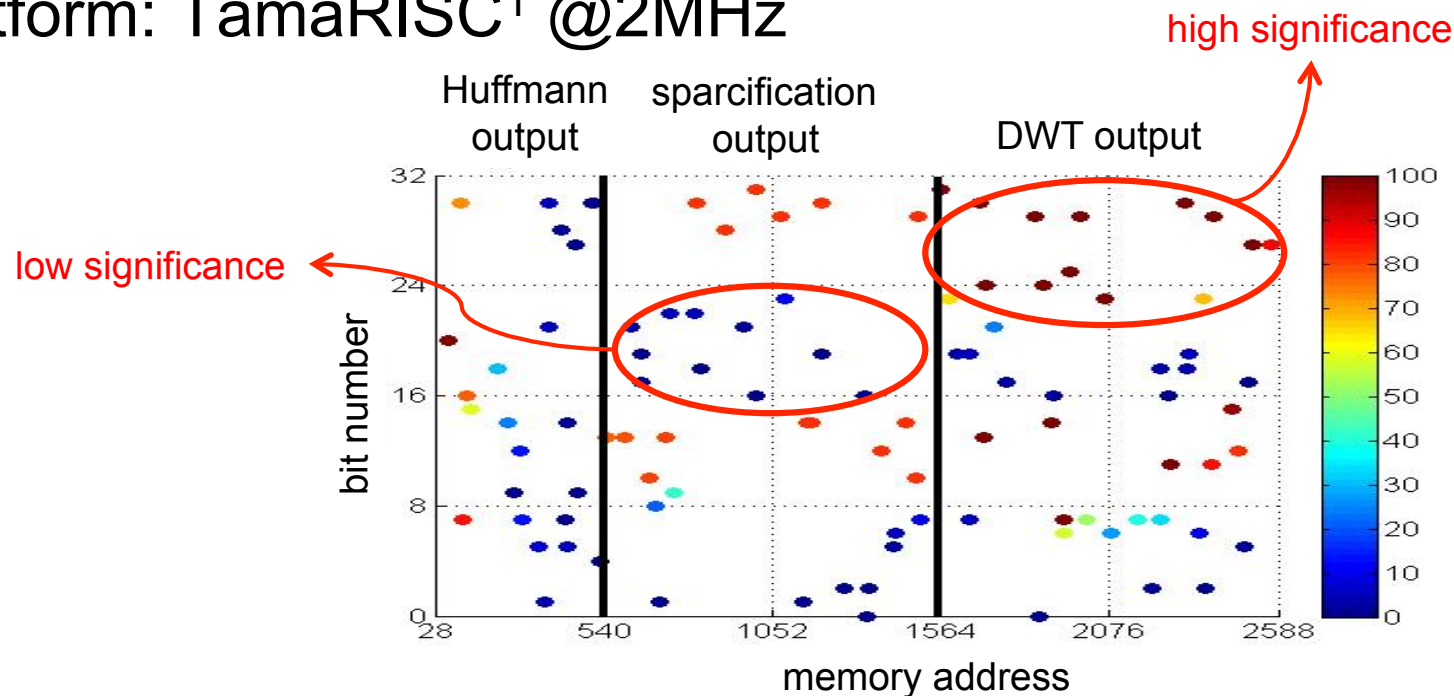


- Linear DWT transformation $\rightarrow y = \Psi^{-1} x$
- Sparcification \rightarrow smallest coefficients are set to zero
- Encoding \rightarrow Huffmann encoding



Sensitivity analysis

- Significance sensitivity analysis
 - Black box approach
 - Inject errors (bit-flips) and observe quality of output
 - ~1000 simulations (50 traces from MIT-Arrhythmia DB, 1 minute each)
- Platform: TamaRISC¹ @2MHz



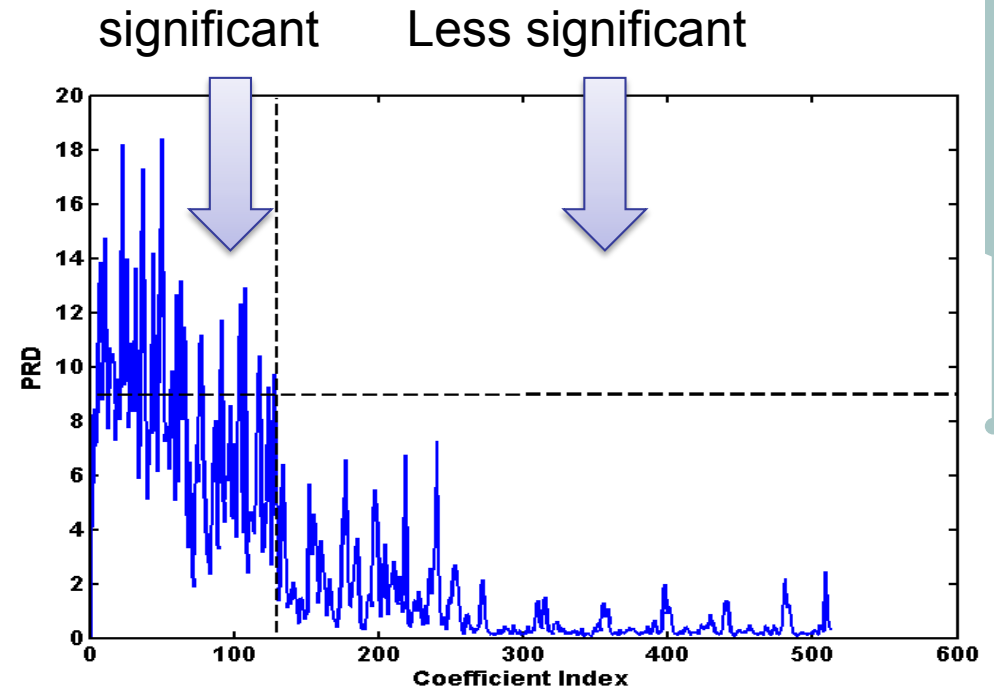
[1] Dogan, A.Y.; Constantin, J.; Atienza, D.; Burg, A.; Benini, L., "Low-power processor architecture exploration for online biomedical signal analysis," Circuits, Devices & Systems, IET , vol.6, no.5, pp.279,286, Sept. 2012

Sensitivity analysis

- Sensitivity metric: Percentage root-mean square difference (PRD)

$$PRD(k) = \frac{\|Y - \bar{Y}_k\|_2}{\|Y\|_2}$$

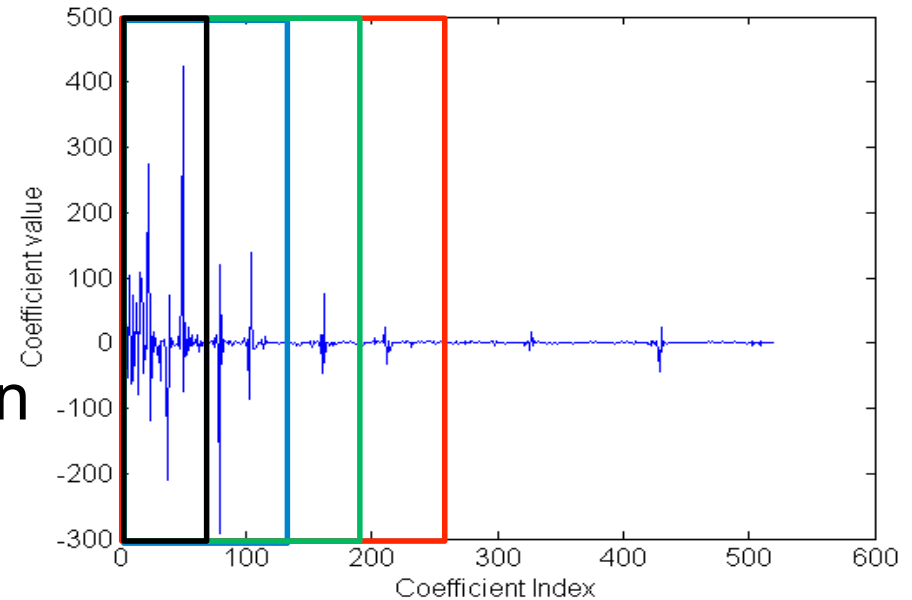
- Sensitivity Evaluation
 - PRD $\leq 9\%$ “good signal”¹
 - PRD $> 9\%$ not suitable signal
- For the DWT step:



[1] Zigel, Y.; Cohen, Arnon; Katz, A., "The weighted diagnostic distortion (WDD) measure for ECG signal compression," Biomedical Engineering, IEEE Transactions on , vol.47, no.11, pp.1422,1430, Nov. 2000

DWT Selected Significance cases

- Less significant data
1-bit parity
 - Only detect errors
 - Remove erroneous data
- Significant data: full protection
4-bits ECC
 - I: 12.5% significant
 - II: 25% significant
 - III: 37.5% significant
 - IV: 50 % significant
- Maximum error occurrence
 - Less significant: total corruption (worst case)



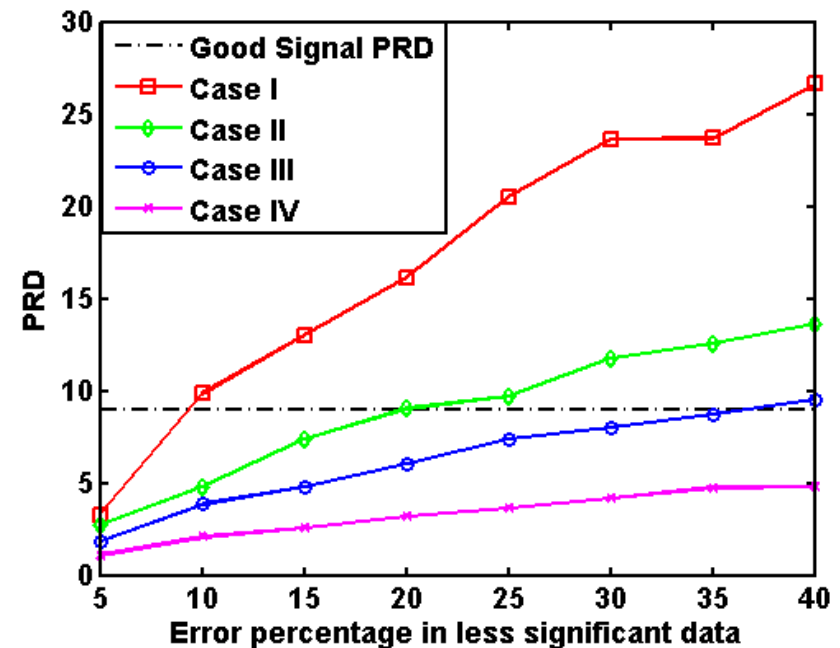
Case	PRD
I	63%
II	32%
III	18%
IV	9%

Experimental Evaluation: PRD

- Error injection
 - Rate: 1e-6 to 1e-9 bit/cycle
 - Effective (practical) maximum error occurrence: 40% of less significant

- Guaranteed good quality signal
 - Case III (37.5% significant)
 - Case IV (50 % significant)

- Time-dependent case-switching
 - Better savings
 - Higher complexity



Experimental evaluation: energy saving

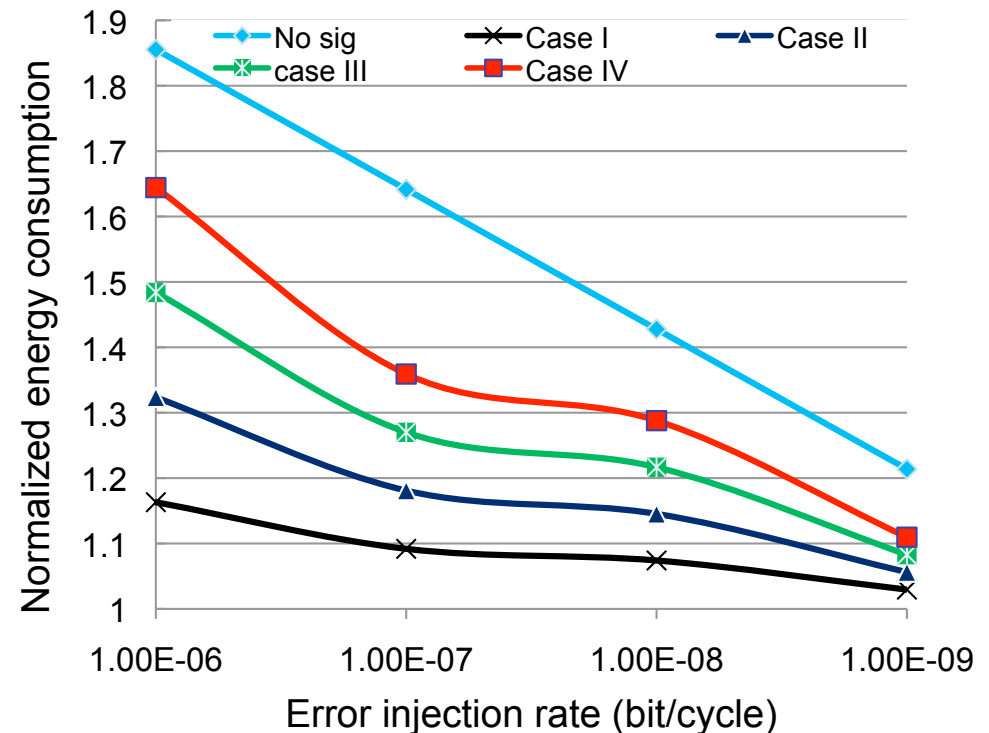
- Protection energy normalized with respect to no protection

- Case I (12.5% significant):**

- 80% less protection energy overhead than full protection (No sig)
- 37% overall energy savings

- Case III (37.5% significant):**

- 43% less protection overhead
- 20% overall energy savings

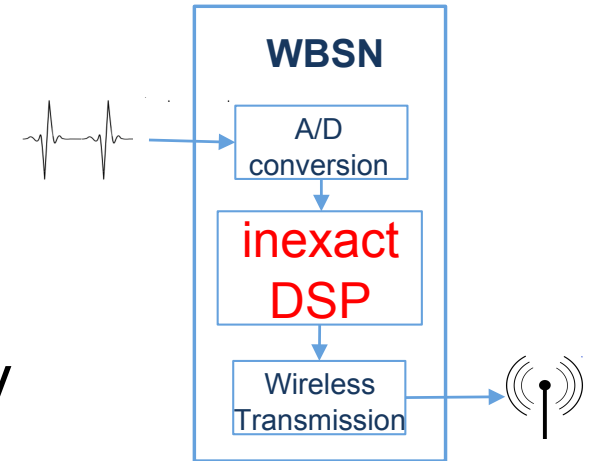


- Also: time savings**

- no correction overhead for non-significant data

Conclusions

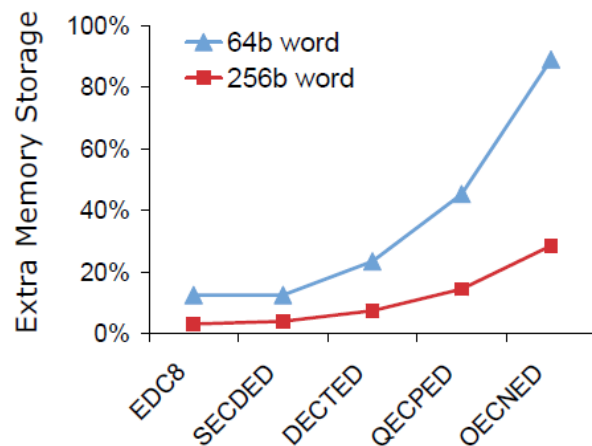
- WBSNs have tight energy constraints
- Noisy inputs, qualitative outputs
- Inexact DSP to increase energy efficiency
 - embedded processing
 - wireless transmissionat both **hardware** and **software** layers
- Complex interactions among analog, digital and wireless transmission stages



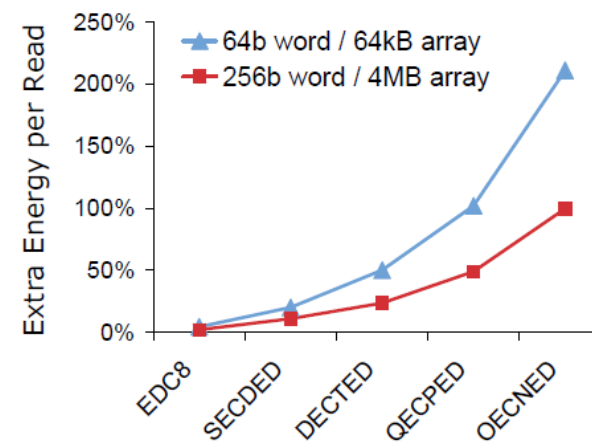
Thank You

Protection mechanisms

- Hardware protection
 - Error correcting codes (ECC)
 - Energy and area inefficient
- Firmware/software protection
 - Backward error correction → Execution time ↑
 - Forward error correction (SW-based ECC) → Complexity ↑



(b) Area overhead of ECC



(c) Energy overhead of ECC

Kim et al. MICRO-40