Towards Smart Wearable Systems in the Internet-of-Things Era

Prof. David Atienza,
École Polytechnique Fédérale de Lausanne (EPFL)

MCSoC-16, Lyon Congress Center, Lyon (FR), Sept. 21-23, 2016

Wearable Systems are.. Many Things

- Many different purposes… And complexities (today more than 3000 products)

[Courtesy: C. Henz]

Wearable Applications: Big Market Expansion

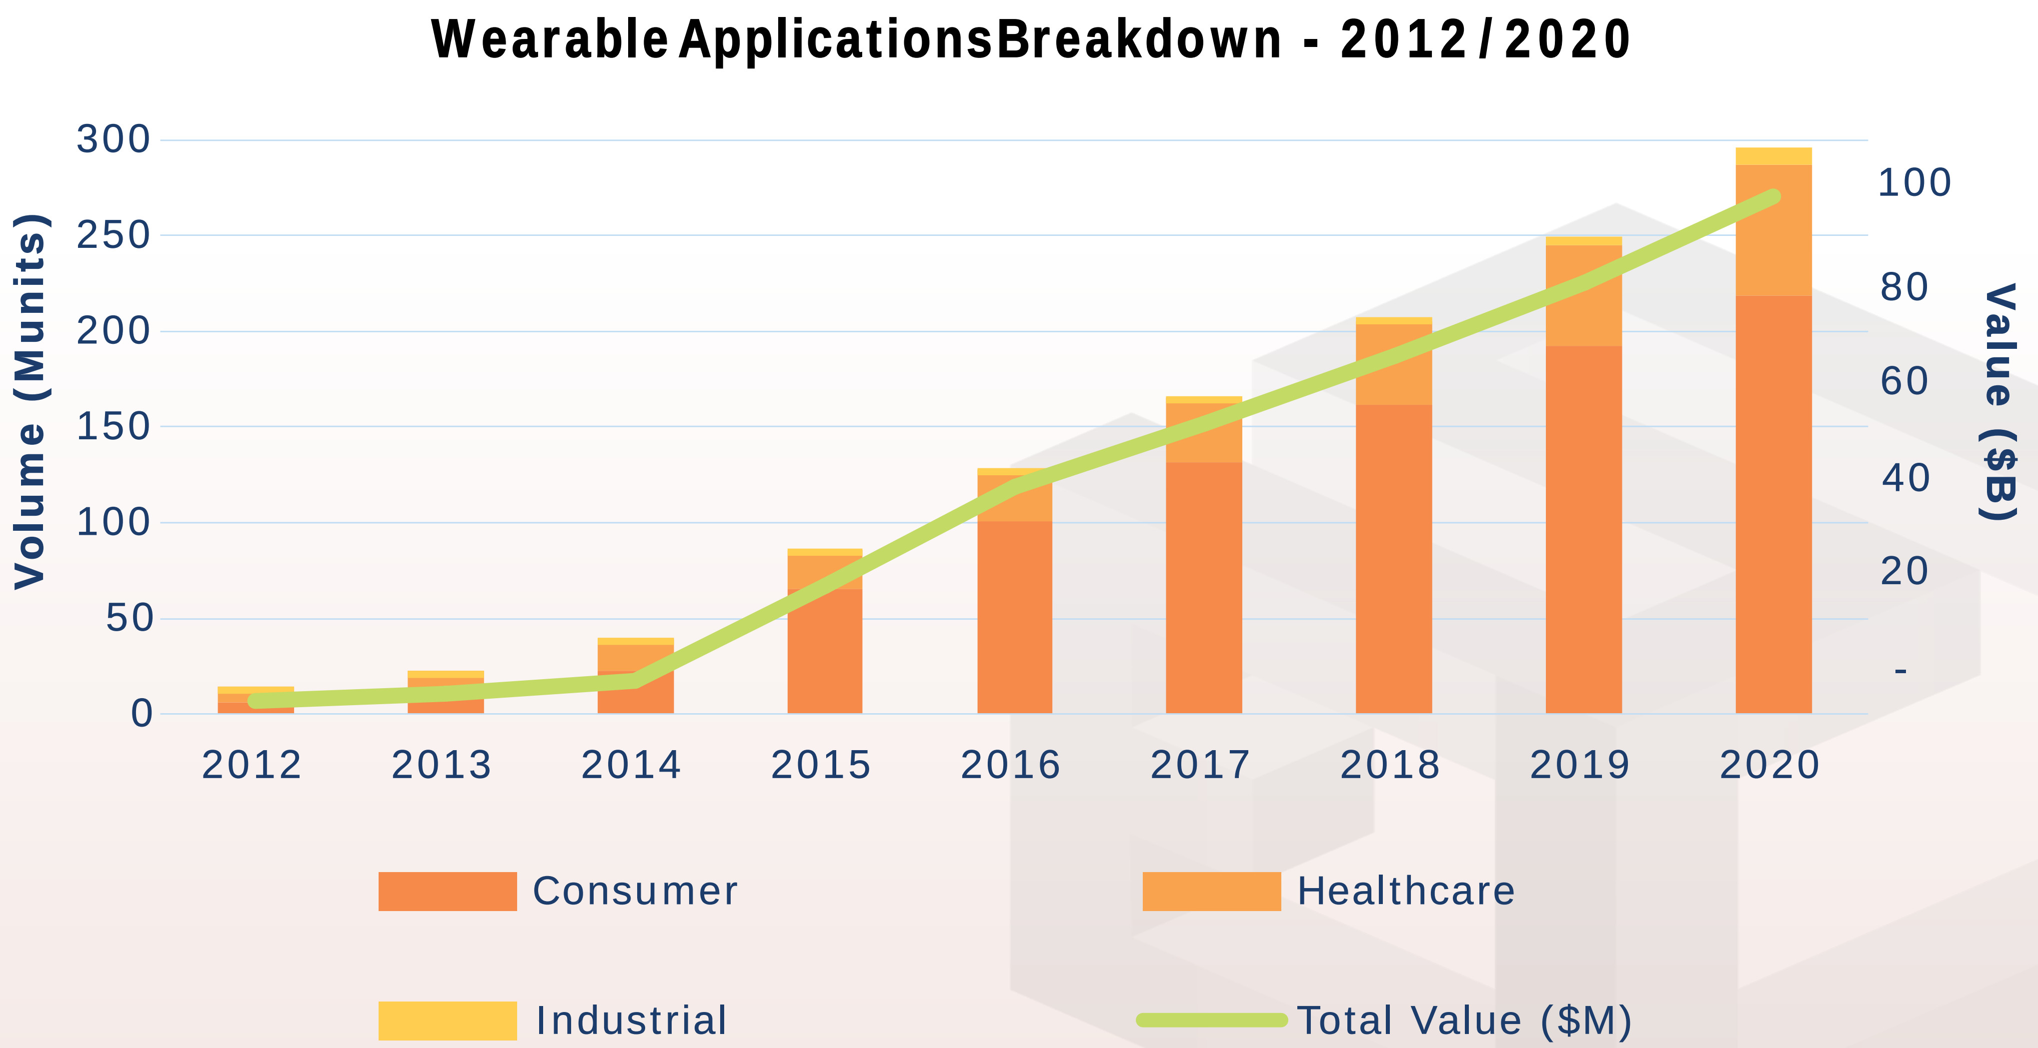

(Yole Développement, July 2015)

Internet-of-Things (IoT)

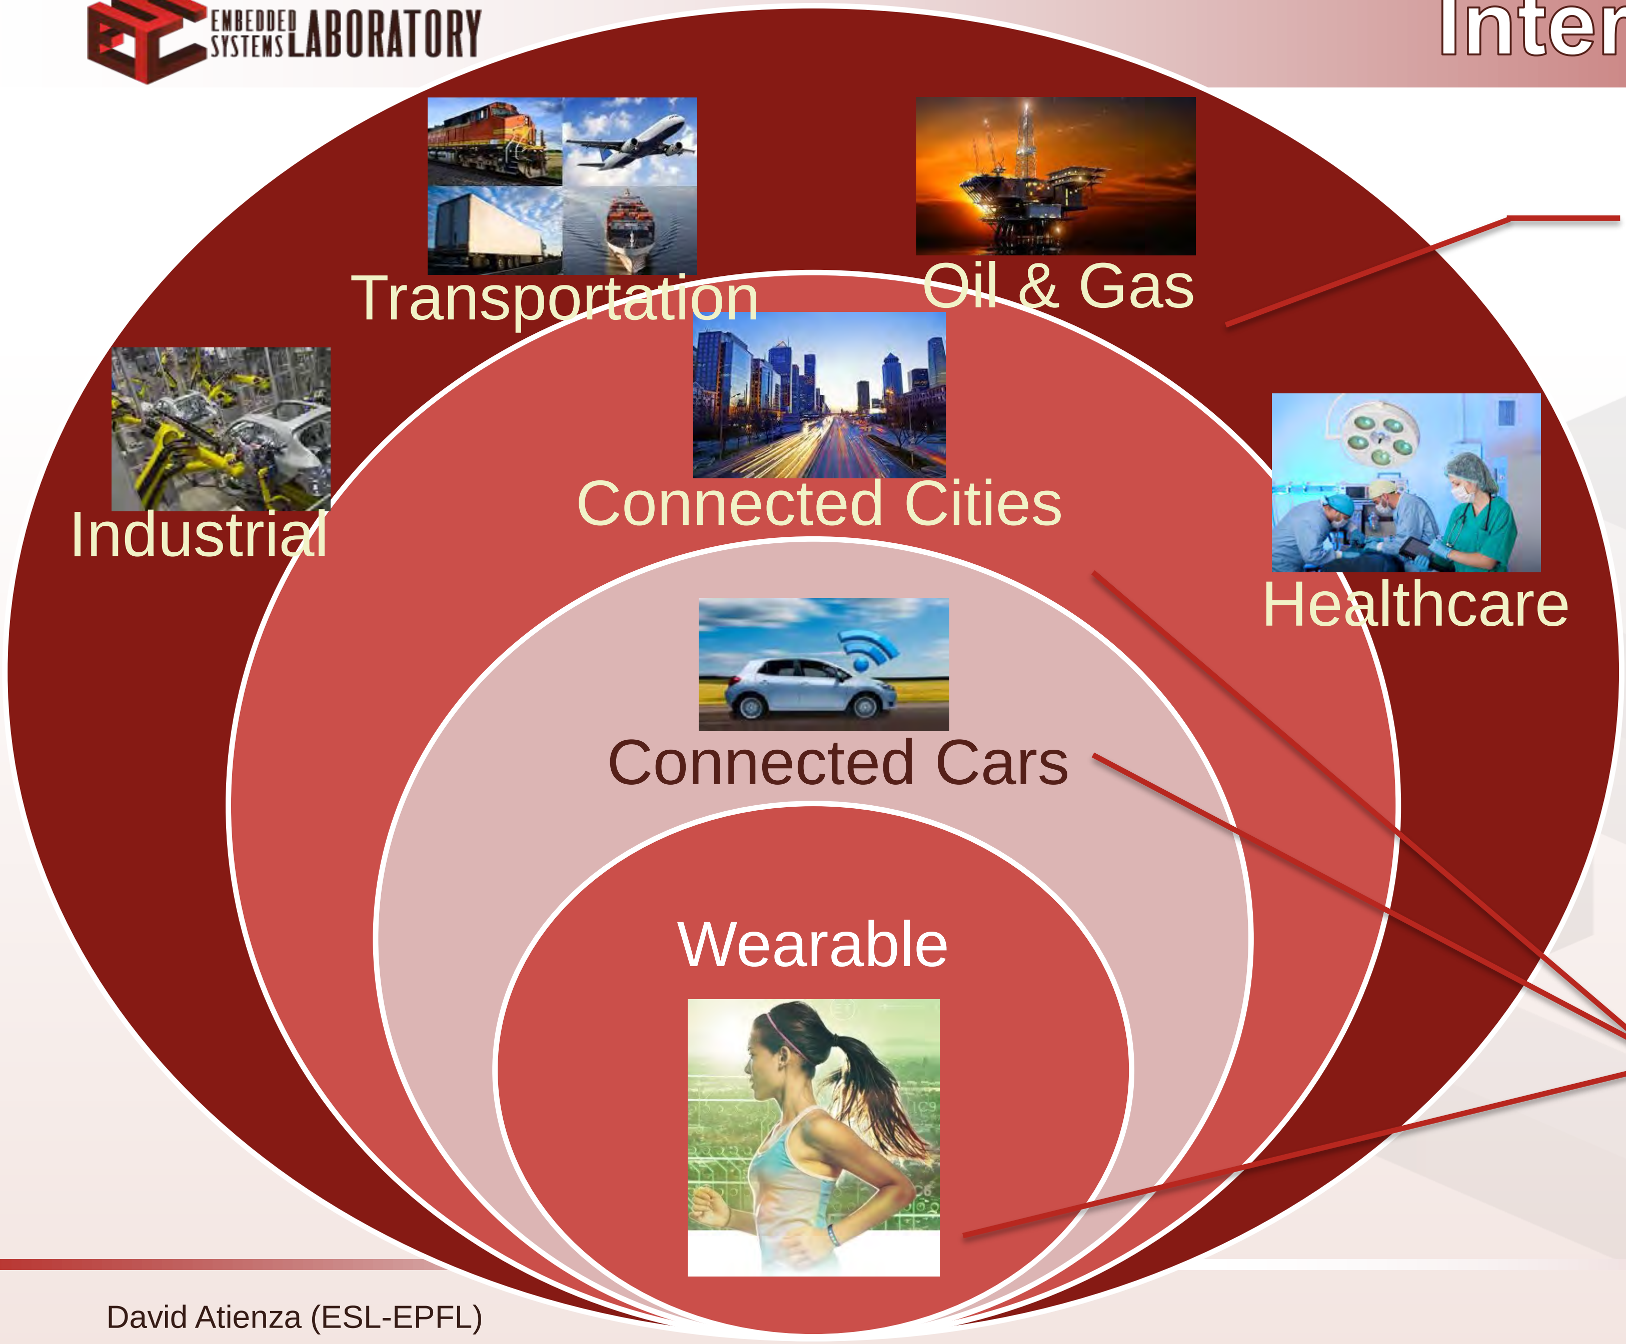

Big Data Analytics Frameworks & Machine Learning Algorithms

Continuous system monitoring

[Source: Goldman Sachs Inv. Res.]

Potential of Internet-of-Things (IoT)

- IoT long-term economic benefits [McKinsey]
  - Remote healthcare: **$11.1Trillion/year saved (1B people)**
  - Efficient energy use: **45TWh/year saved in EU (4M houses)**
  - Business-to-business services: **70% added value**

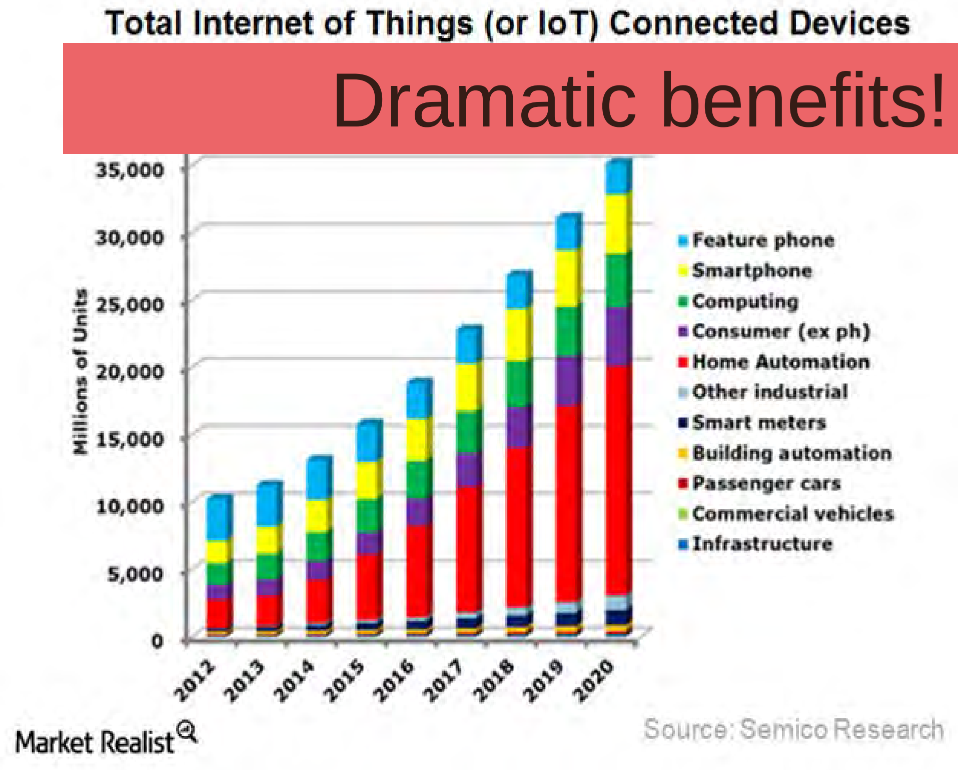

Dramatic benefits! But will this really work well?

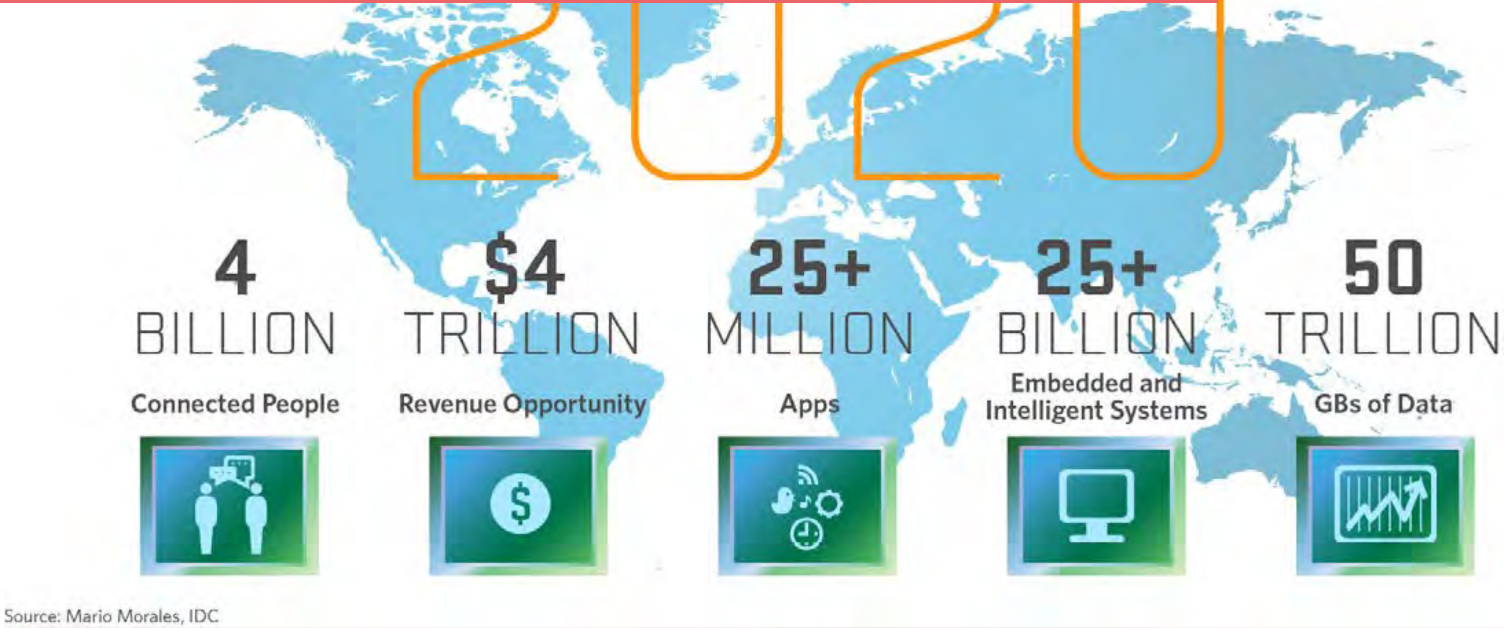

Wearable Systems in IoT Era to Rescue Our Healthcare

- Burden of disease shifted in recent years
  - Disorders with behavioral causes are key
  - Expected to be 75% of GDP by 2030 [McKinsey

- Two-fold paradigm shift in health delivery

Symptom-based	➔	**Preventive healthcare**
Hospital-centered	➔	**Person-centered**

- Cardiovascular monitoring is key today…

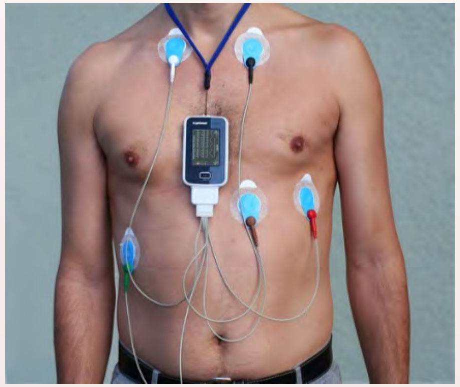

ECG Holter data logger (clinical practice)

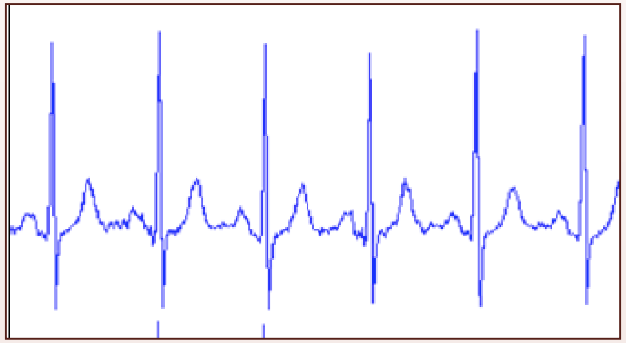

Resting Electrocardiogram (ECG)

Wearables in IoT era will relay information to the cloud and healthcare providers

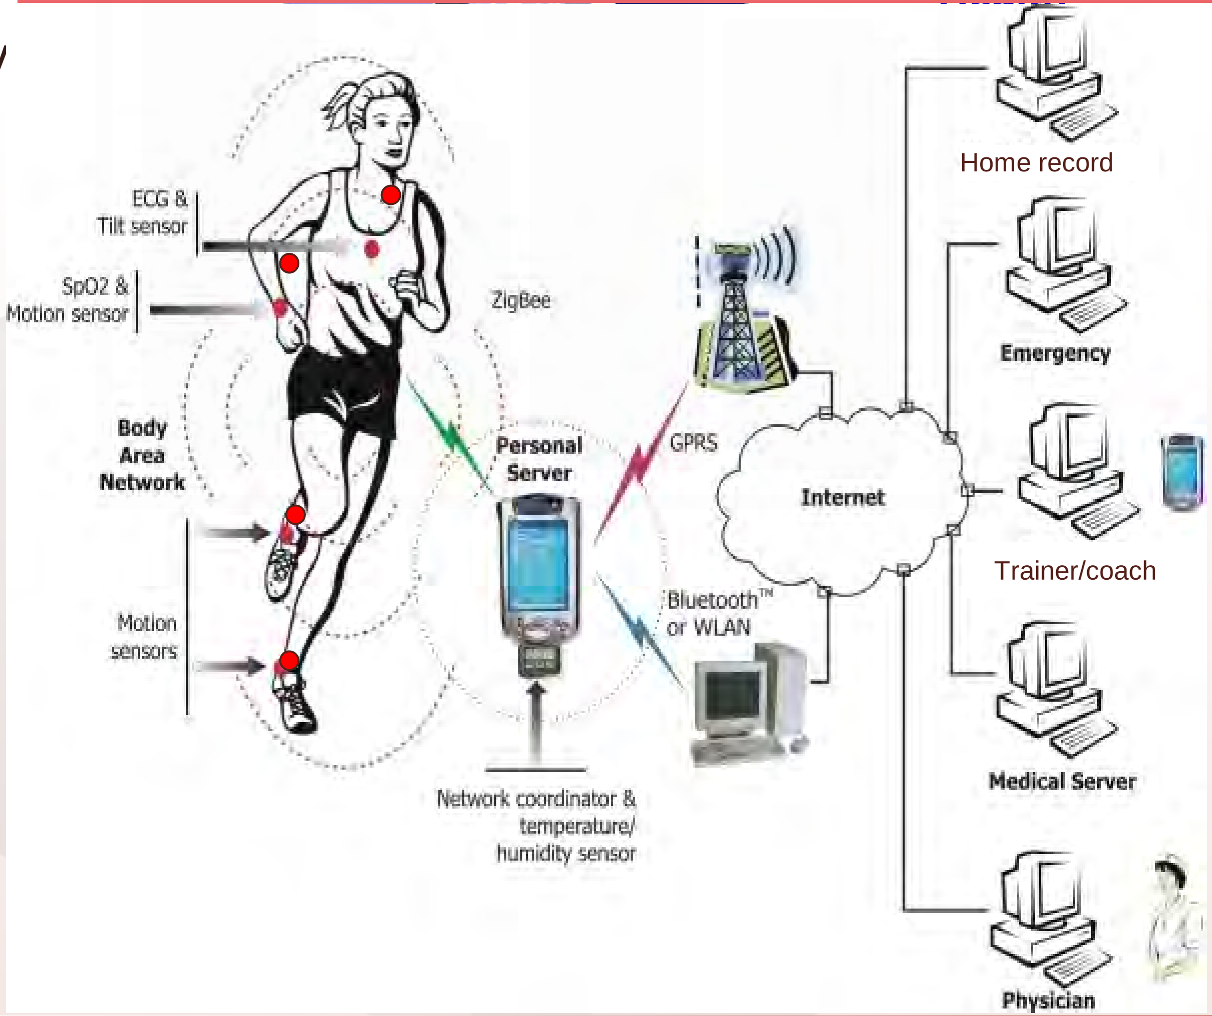

State-of-the-Art Wearables Today

Shimmer (shimmerresearch.com, 2014)

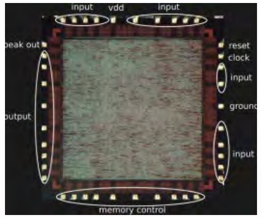

Heart Rate Monitoring (Massagram, 2010)

Corventis's PiiX (Corventis cardiac systems)

Toumaz's Sensium Life (Wong,2012)

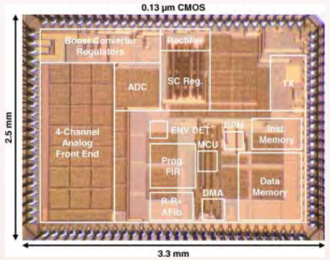

Zhang, 2012

IMEC cardiac patch (Yazicioglu,2009)

Holst Centre (Masse, 2014)

Apple Watch (Apple Inc, 2015)

Raw bio-signal transfer or after simple tasks (heart rate monitoring, basic filtering) on single-input biosignals

State-of-the-art Wearable Node: the Shimmer™ platform

- **TI MSP430 microcontroller**
  - 16-bit, 8MHz, 10KB RAM, 48KB Flash
  - ADC converters, DMA, HW multiplier
- **CC2420 radio**
  - 250 Kbps, ZigBee compliant
- **Sensors**
  - 3-channel ECG
  - Accerelometers and gyroscopes

- CONSTRAINTS:
  - No floating point operation
  - No hardware division
  - Limited memory
  - Limited computing power
  - Limited autonomy (rechargeable Li-polymer battery of 250 mAh)

Long-Lived Wearable ECG Monitors Require Major Design Shift

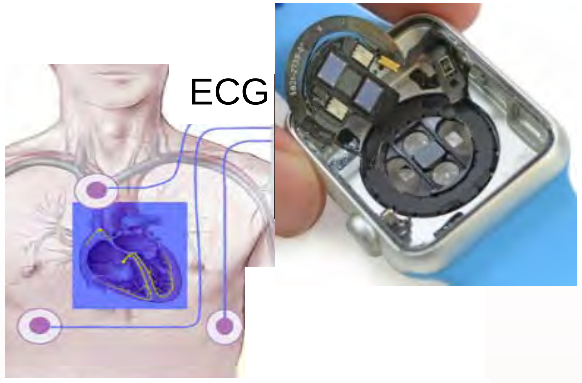

1. Can we reduce the data sensing/sampling cost and the amount of streamed data?
2. Can we embed automated analysis without compromising the system lifetime?

Under stringent processing and memory constraints… **Power!**

- This wireless 1-lead ECG streaming monitor lasts 134.6 h (2011)
  - Current wearable technology it lasts 172.5 h (2016)

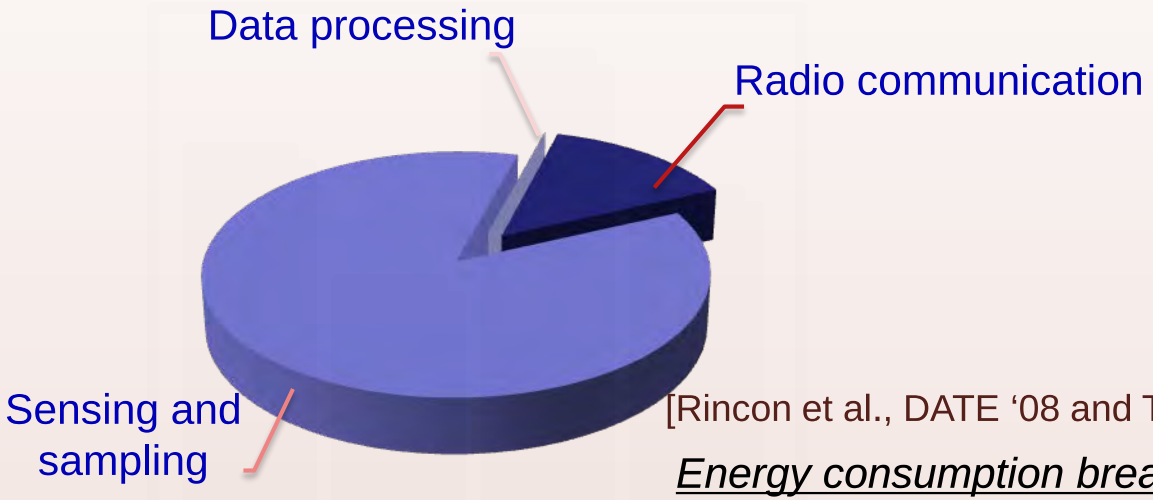

[Rincon et al., DATE '08 and TITB '11]

Energy consumption breakdown

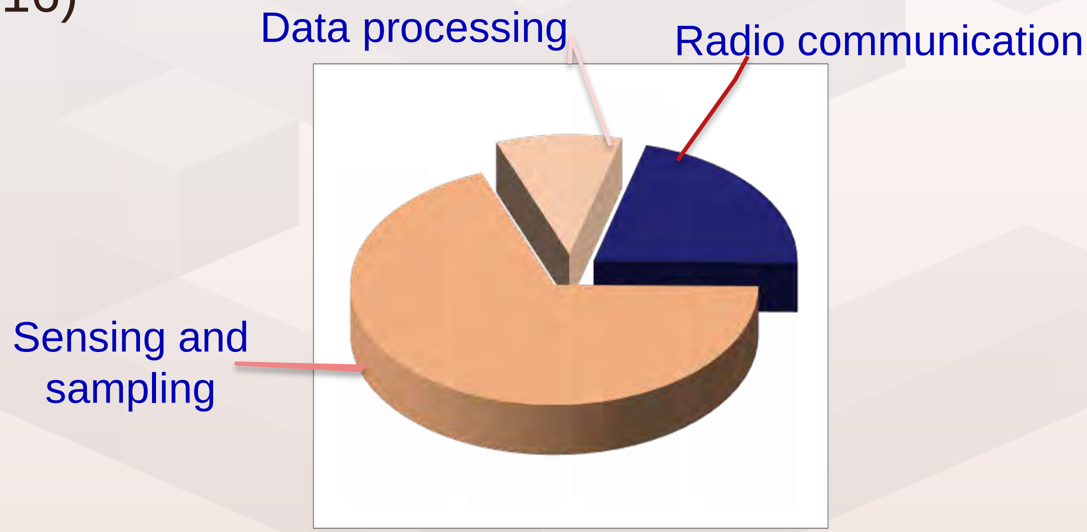

IoT on Top of the Wave… How Long?

- Diffusion of innovations: very high expectations since 2014
  - But no real progress apparently (or very slow at least), too complex or just impossible?

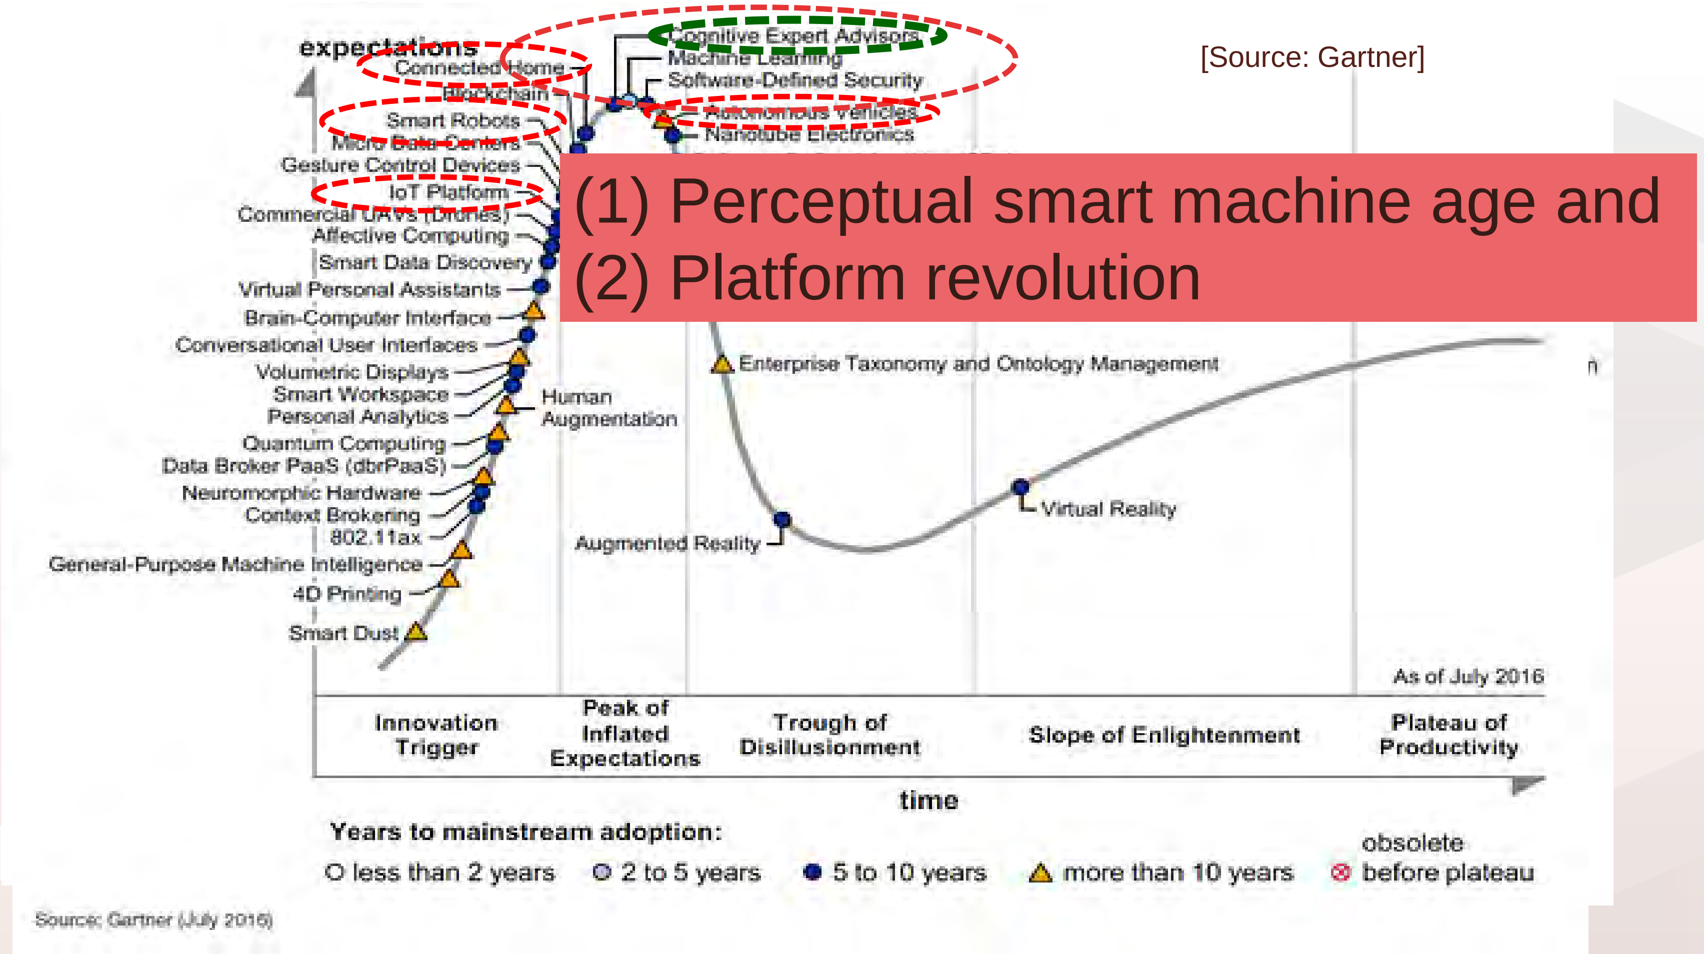

IoT Unaffordable in Healthcare With Current IoT Approach

- Cheap ... ns
  - Sen
  - Ban ... ars
  - Netw
- But no
  - **Con**
  - **Poo**
- **Let's s**

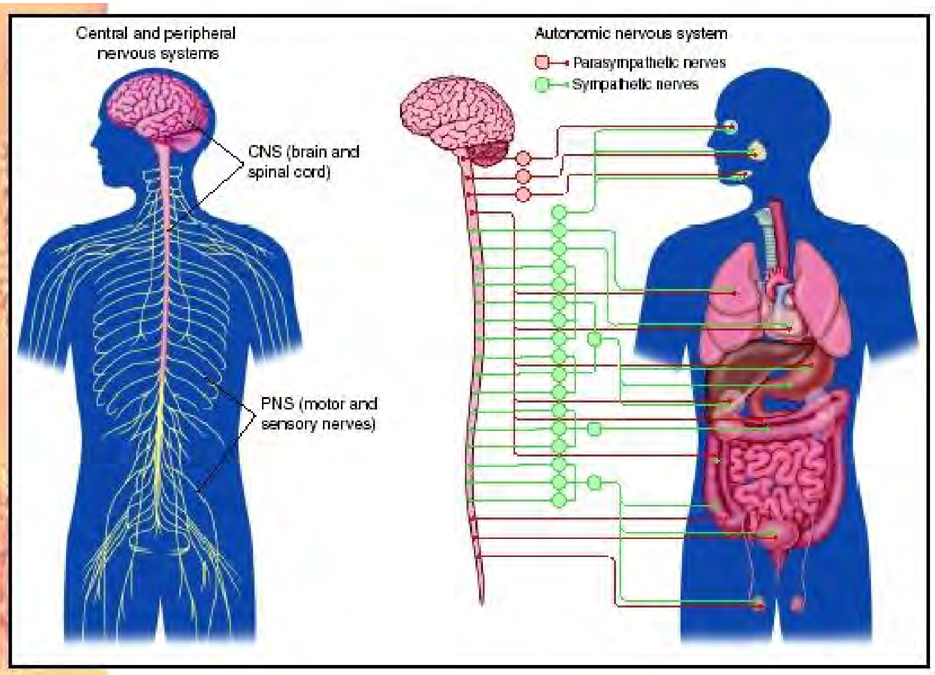

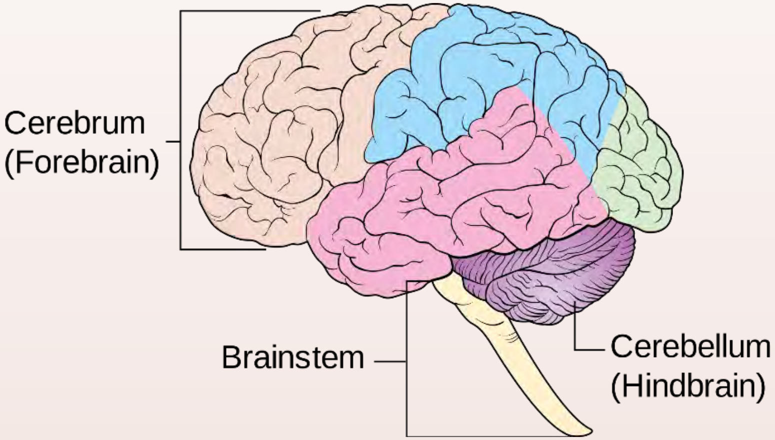

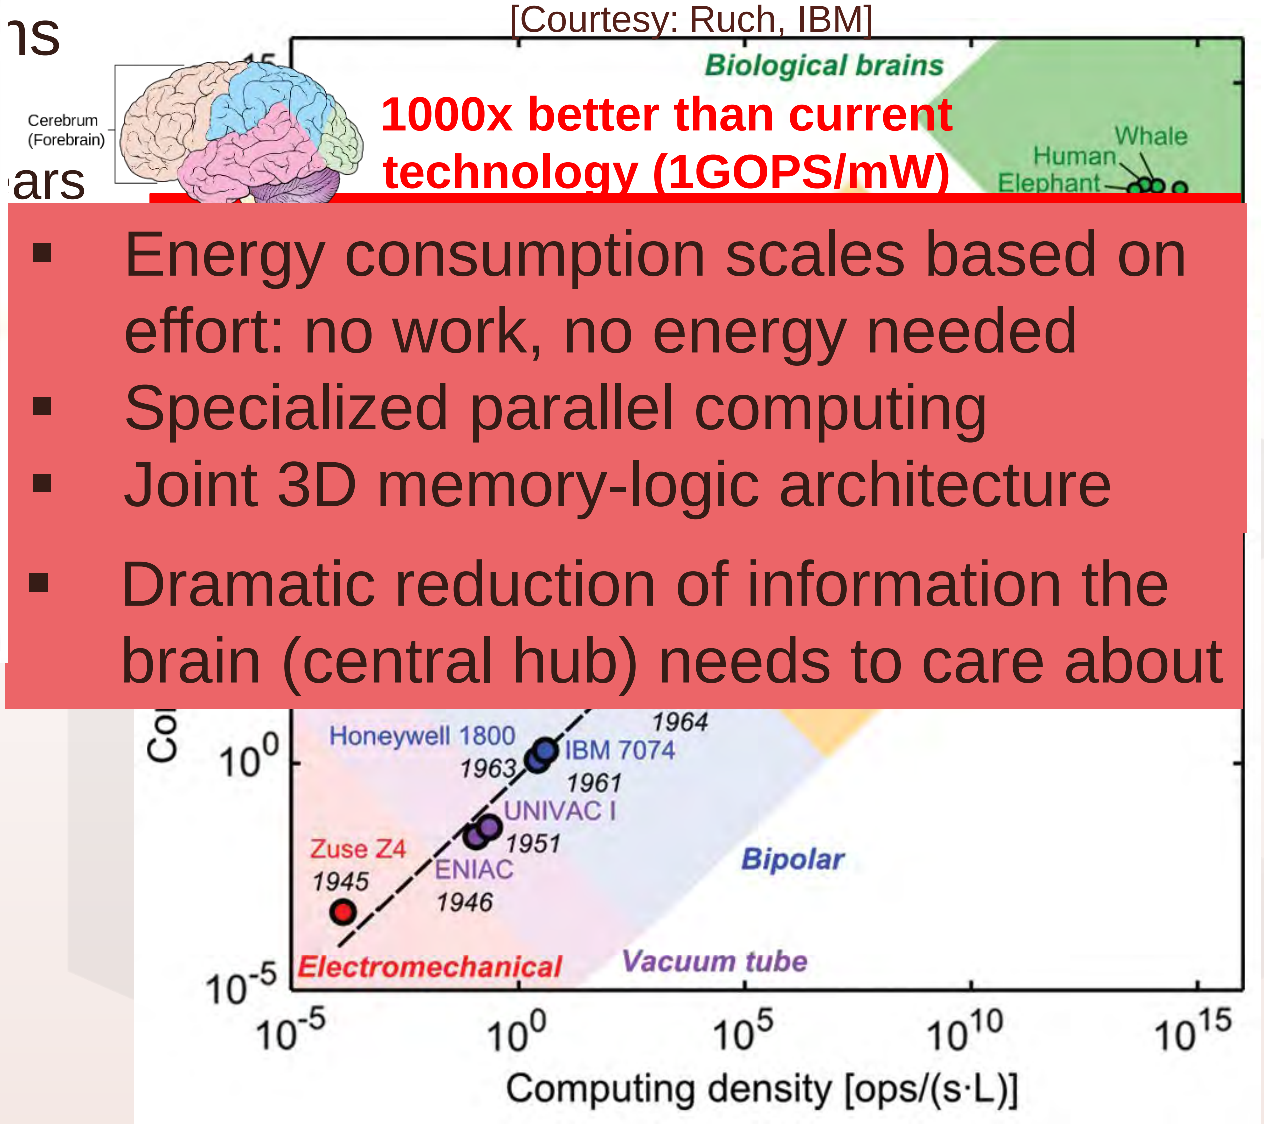

- Energy consumption scales based on effort: no work, no energy needed
- Specialized parallel computing
- Joint 3D memory-logic architecture
- Dramatic reduction of information the brain (central hub) needs to care about

Smart Wearable Systems for Autonomous Arrhythmia Monitoring

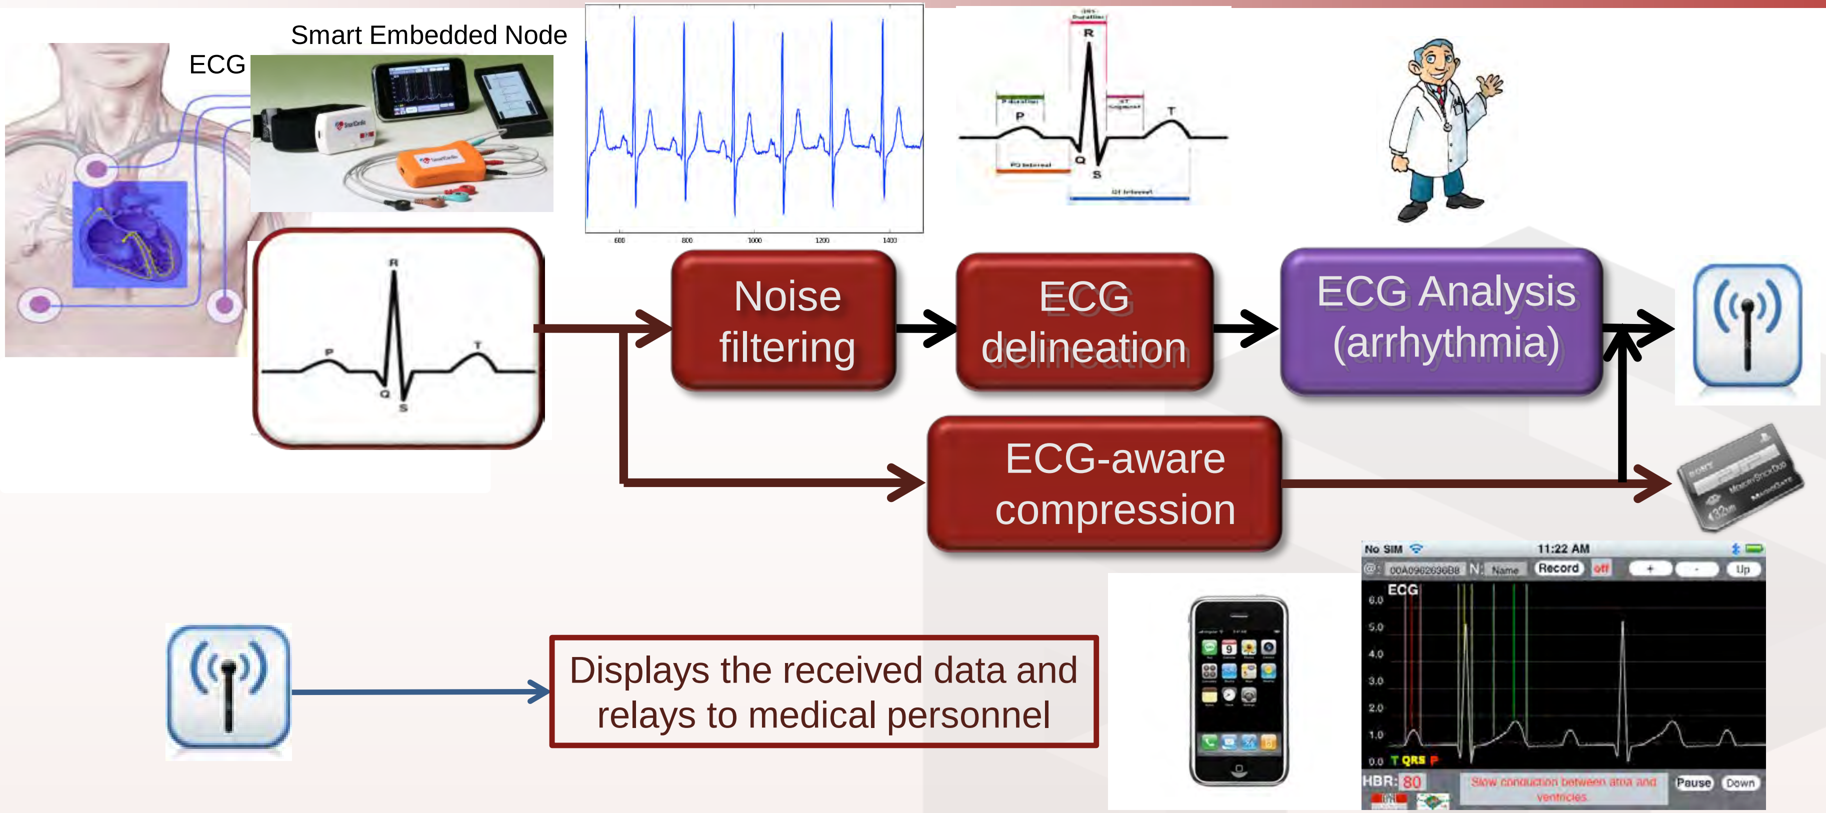

Software: wearable systems can implement multi-lead ECG analysis

- Filtering: Low-complexity methods using integer computing **(real-life tests on measured points)**
- Delineation: Multi-lead ECG arrhythmia analysis in real-time (**doctor support for quality loss**)
- Communication and storage: bio-signal based compressive sensing (**only 30% of data kept)**

Personal Arrhythmia Smart Wearable Device

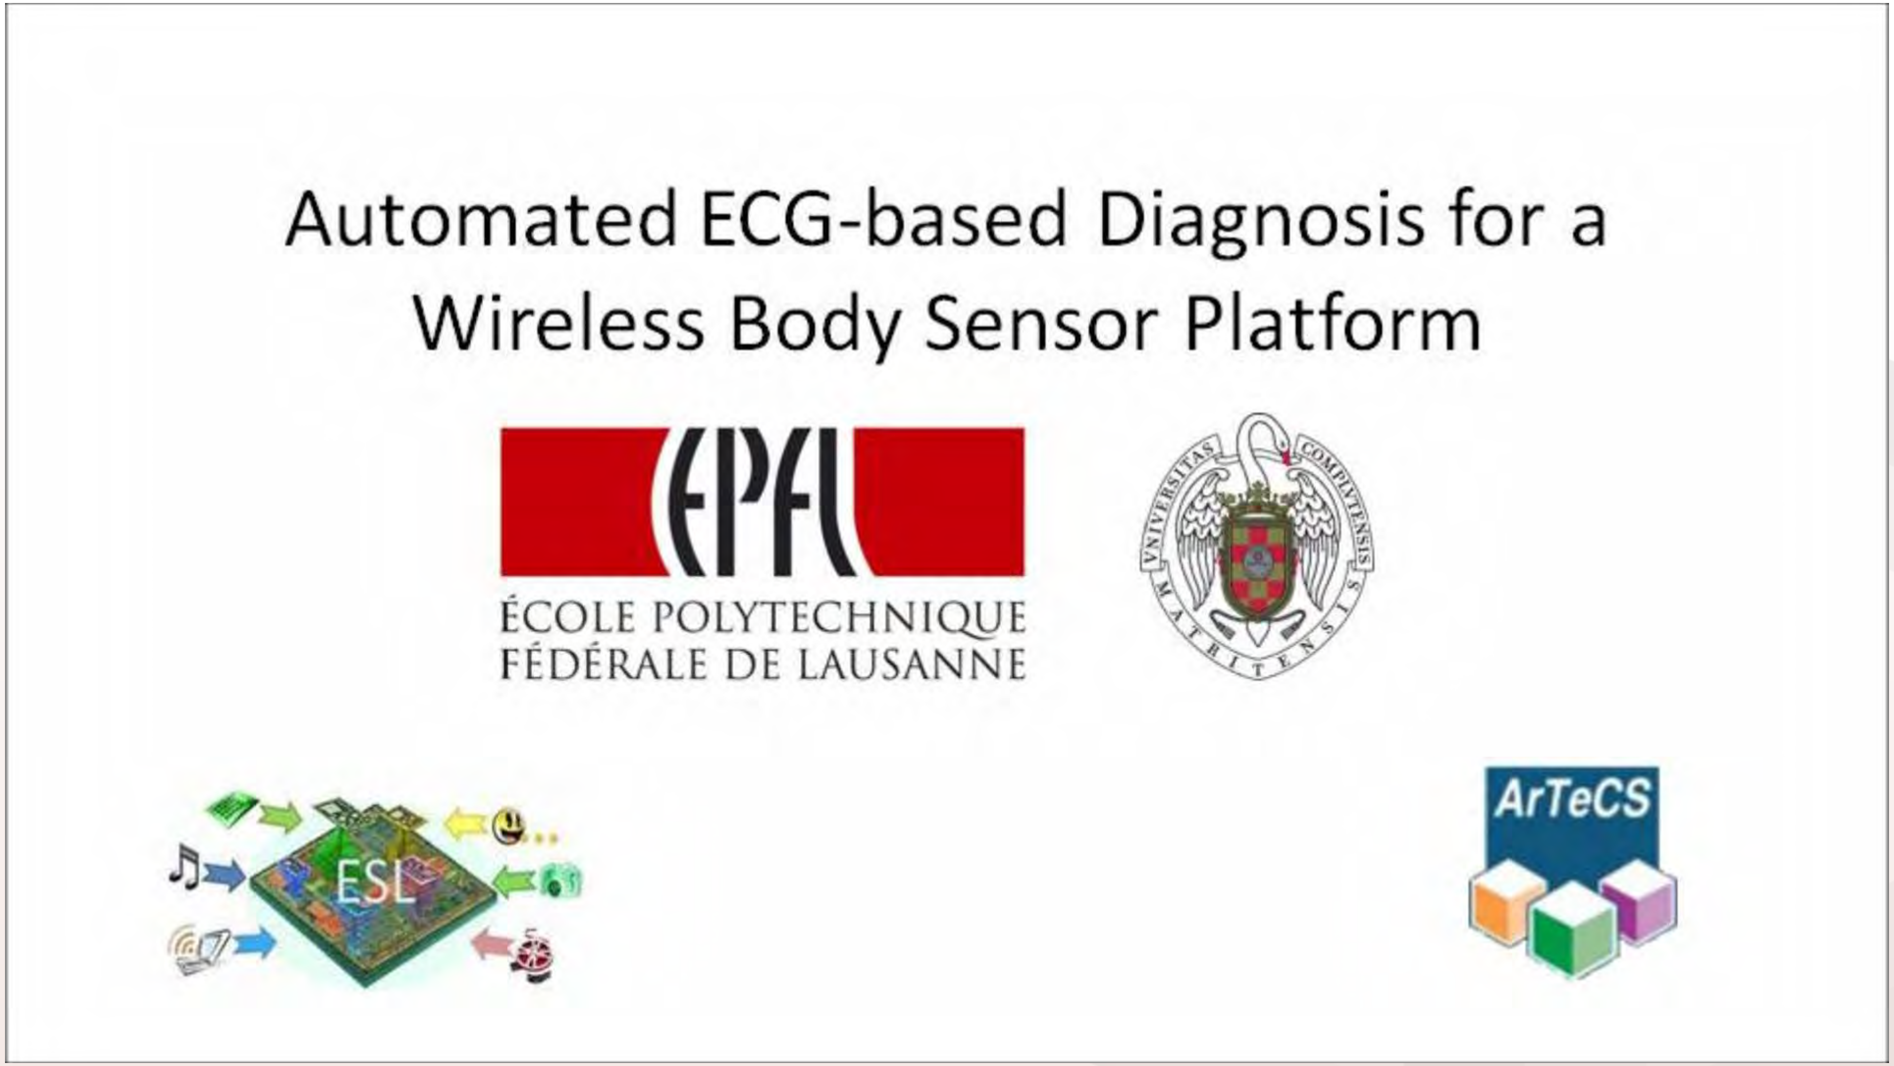

See video at: http://esl.epfl.ch/cms/lang/en/pid/46016

Results on the Shimmer as Smart Wearable System

- Real-time delineation demands limited requirements after careful algorithm optimization (computational load and memory footprint)

Algorithm	RAM usage	Buffers length	Execution time
Single-lead WT delineator	6.8 kBytes	512 elements	5%
Multi-lead WT delineator (morphological filter of baseline removal)	5.5 kBytes	256 elements	30.5% total (23% filtering, 2.5% multi-lead merging, 5% delineation)

Execution of complex automatic ECG processing algorithms is possible

Small on-chip memory (10 kB) is the current limiting factor

Advanced on-chip processing gives real-time information about heart health with **no impact on node lifetime**: more than 139 hours

- Using CS it is sufficient to collect *M* (<<*N*) linear random measurements (samples)

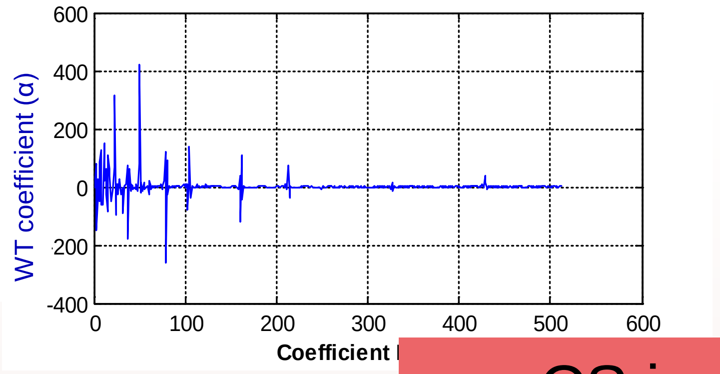

Measurement/Sensing matrix (Gaussian random matrix)

$$y_{M\times 1} = \Phi_{M\times N} \cdot x_{N\times 1}$$

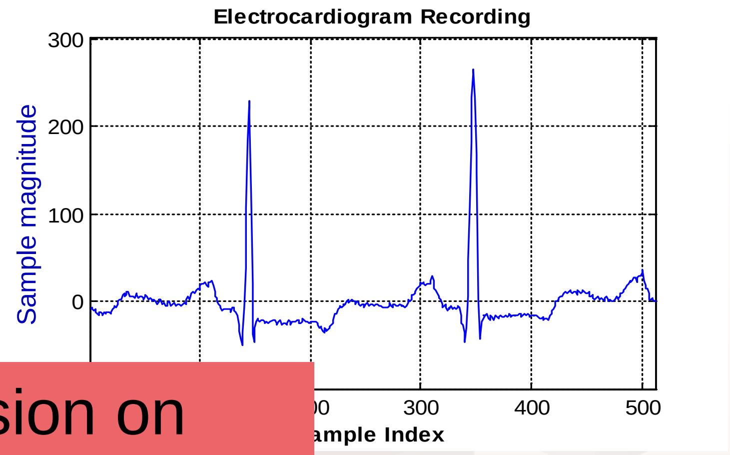

CS is attractive for real-time ECG compression on resource-constrained WBSN, but what about **biosignal degradation** due to CS reconstruction (in real-time)?

- Then, α can be recovered by solving the convex optimization problem:

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

CS-Based ECG Sensor (only 30% of ECG data kept)

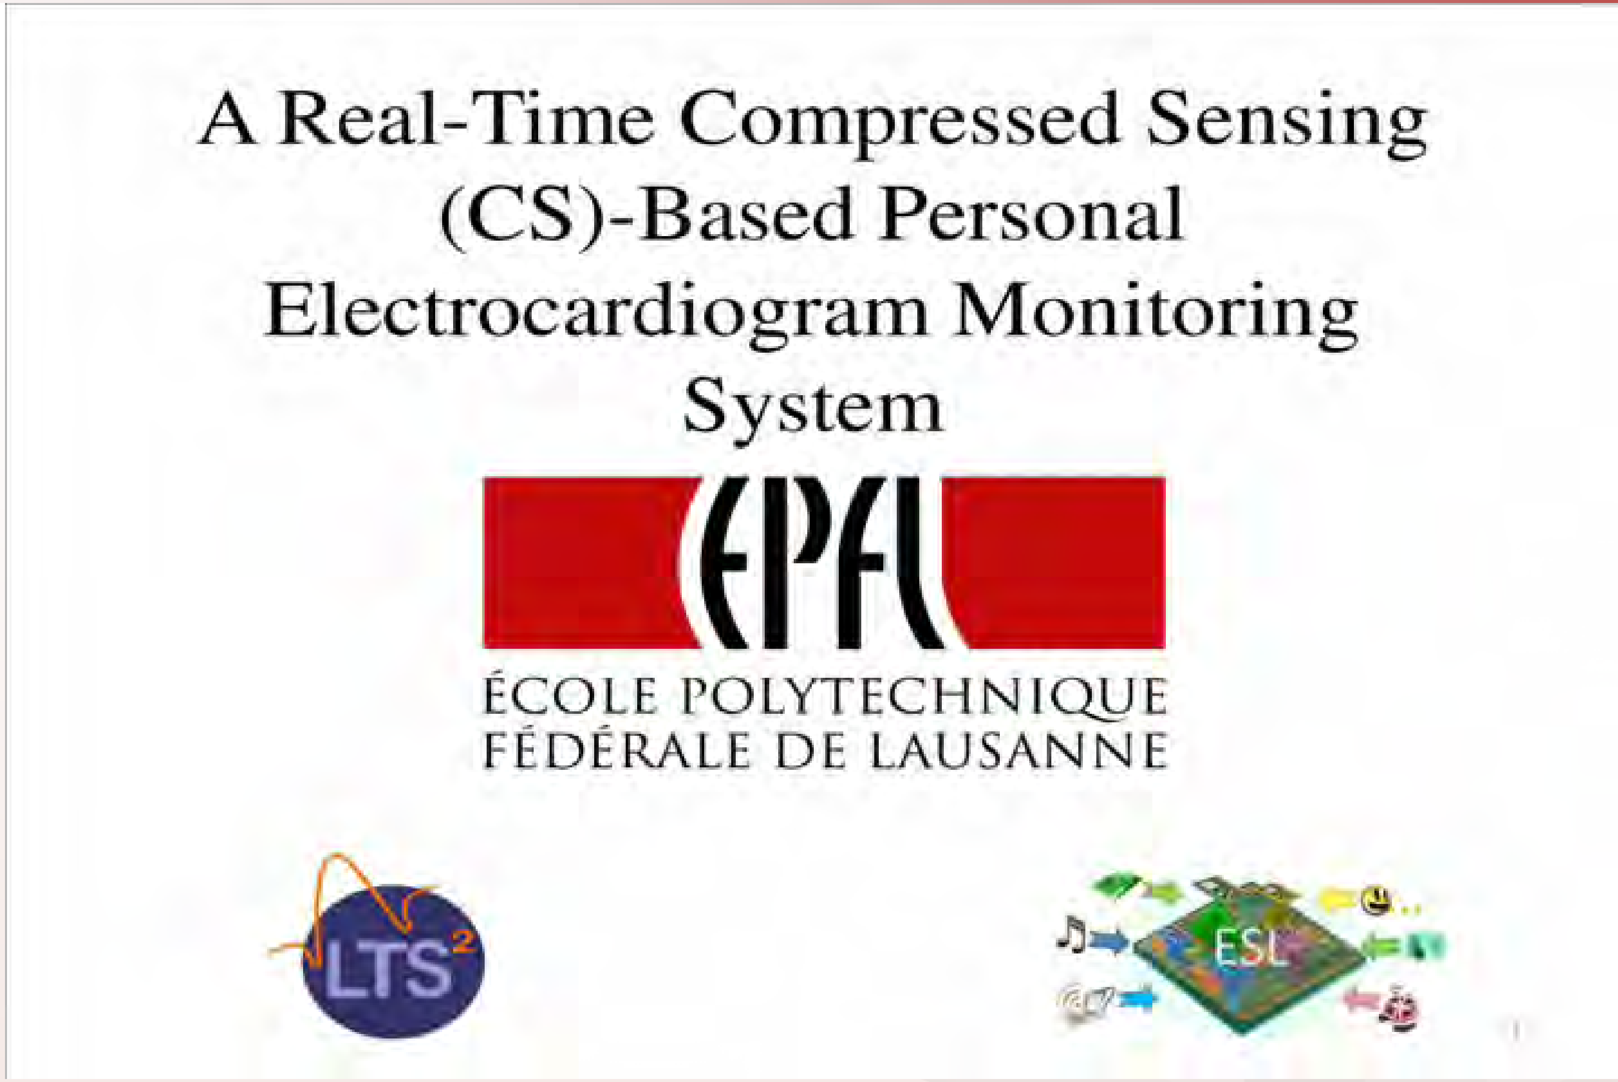

See video at: http://esl.epfl.ch/page-42817.html

CS Gives 23-Fold Reduction in Execution Time, But Limited Lifetime extension

Code execution time

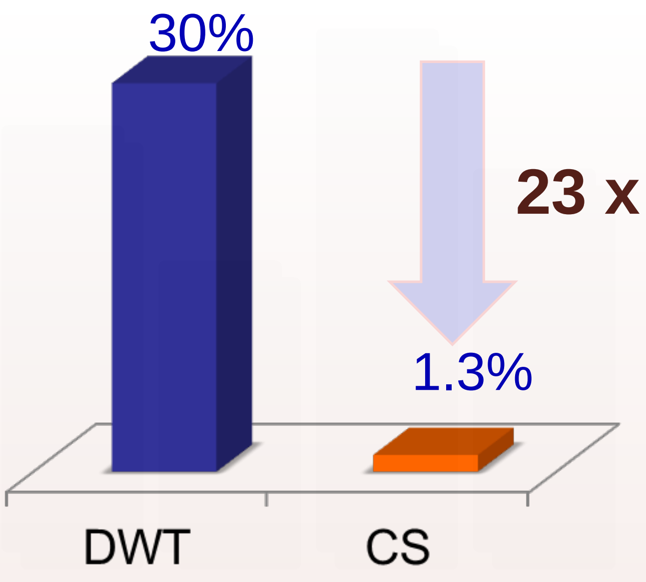

Node lifetime

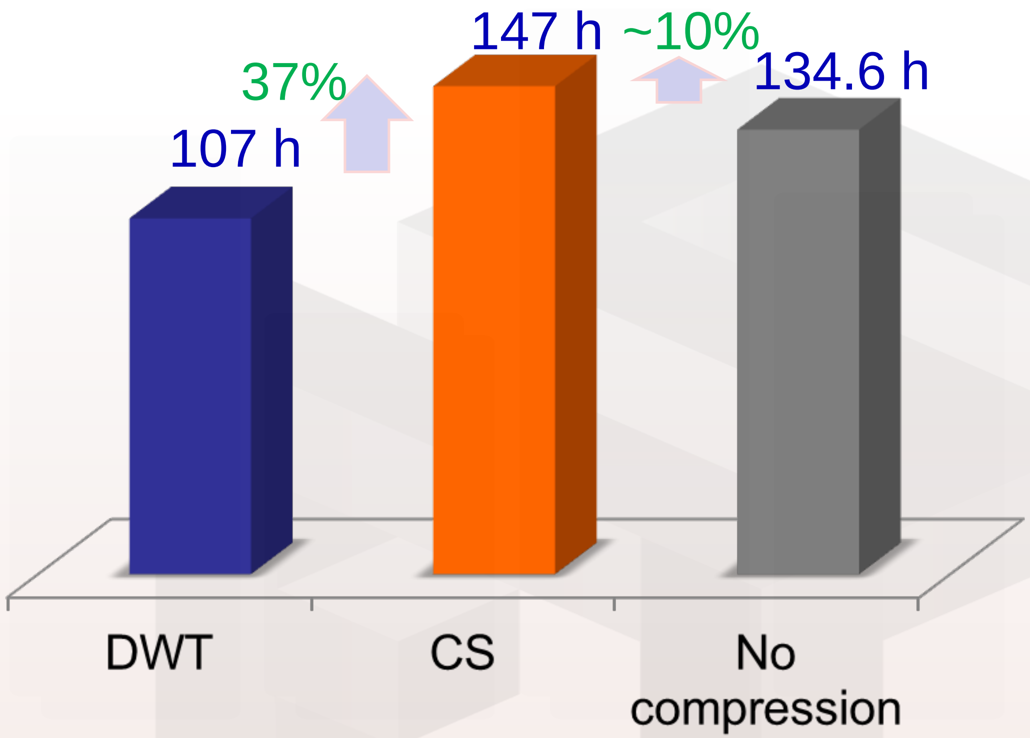

Limited gains because the used generic microcontroller is not optimized for ultra-low-power DSP and CS-based **operations in biological signals**

Energy Reduced Slightly, And Bottleneck Shifted in Smart Wearables

- Initial streaming design
  - **Sensing** dominates (90%)

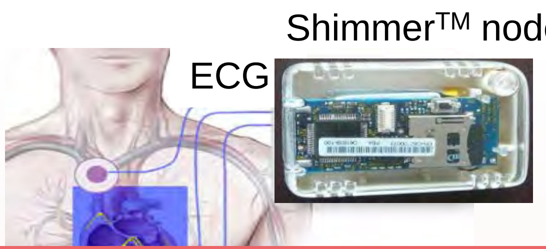

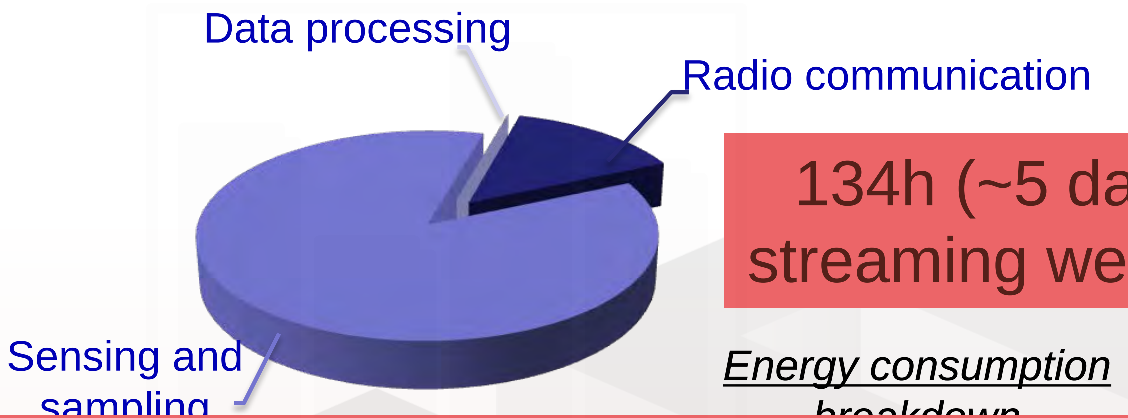

Energy consumption

134h (~5 days) in streaming wearables

Limited gains because the used generic microcontroller is not optimized for ultra-low-power DSP and CS-based **operations in biological signals**

- Smart WBSN M-IoT design
  - **Processing** dominates (70%)
  - **Transmission** still important (23%)

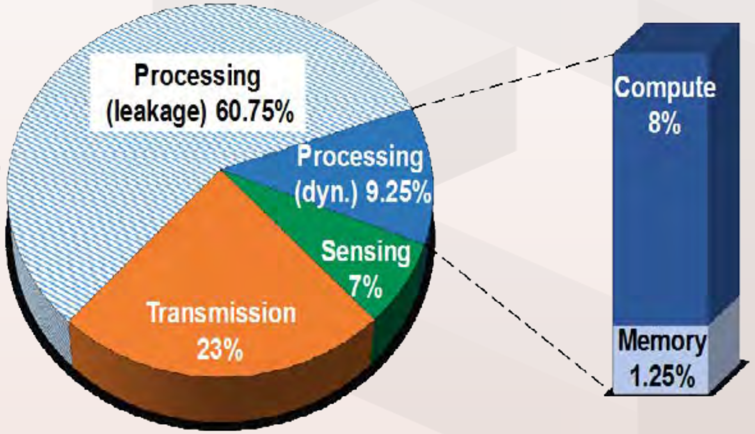

175h (~7 days) in smart wearables

FIRAT: Low-Power Core for Wearables

- Lightweight architecture
  - 3-stages pipeline, 16-bit Harvard RISC (PIC24)
  - 24-bit instruction size
  - Appropriate for single word instructions
- Optimized instruction execution
  - Mostly single word instructions
  - Mostly single cycle execution, low clock
  - Partial data bypass
- Specialized instructions (low clock)
  - Multi-bit shift in one cycle
  - Single-cycle multiplication
  - Multiple-cycle division

[Dogan et al., DATE 2011, DATE 2012]

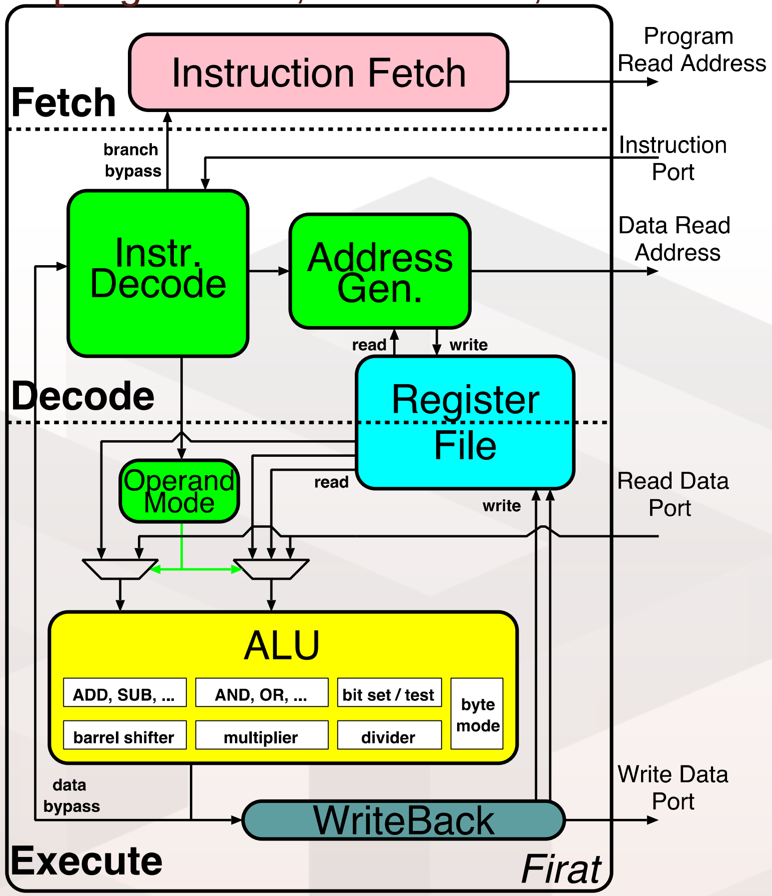

Firat v1.0 (umcL 130nm)

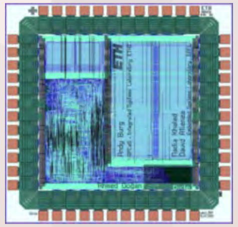

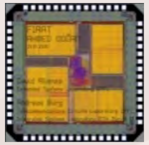

Firat (umcL 90nm)

FIRAT: Experimental results (Wearables for ECG)

	Number of Clock Cycles(*)		
	FIRAT	MSP430	PIC24
Filtering	349 K	658 K	430 K
Compression	114 K	800 K	121 K

(*) 1-package compression (512 samples)

FIRAT vs. MSP430: up to 85% cycle count reduction

FIRAT vs. PIC24: Slightly faster due to enhanced data bypass

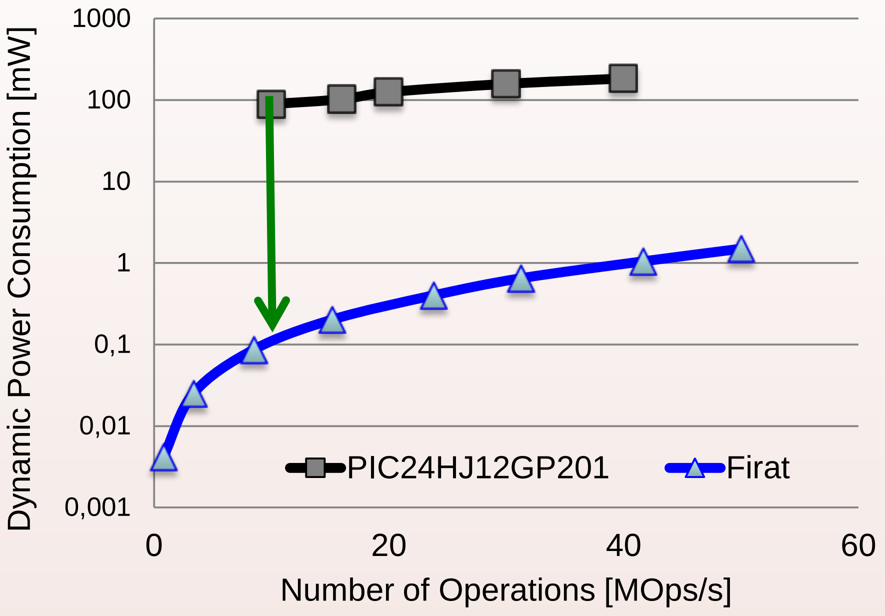

Exploited voltage scaling: Up to 890x power saving with respect to commercial PIC24 chip

So finally we get 2x system lifetime… **And Smart Wearables!**

Smart Wearable Sensors Valuable For Medical Community: Start-Up Case

- Non-intrusive, light and can reduce visits of patients by 50-60% (4-week test)

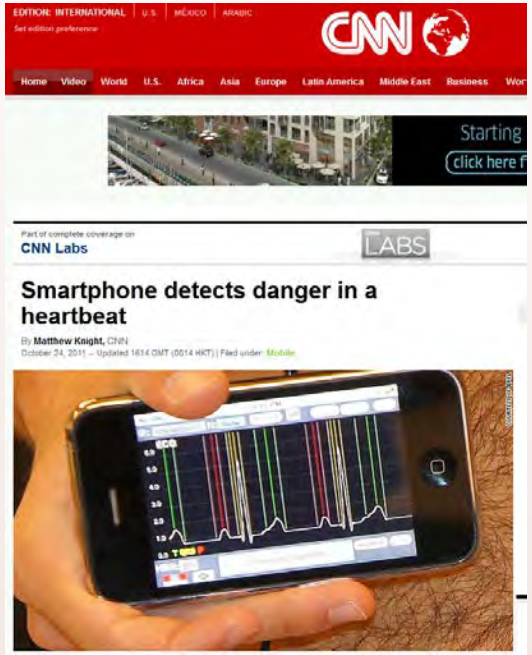

Date: 19.10.2011

Blick

Ein SMS vom Herz

Lausanne – Diagnose: Herzinfarkt. Der häufigsten Todesursache der Welt wird der Kampf angesagt, und zwar mit Schweizer Technik. Forscher der ETH Lausanne haben ein Gerät entwickelt, das den Herzrhythmus konstant überwachen kann. Falls eine Rhythmusstörung auftritt, **sendet das Gerät an Patient und Arzt per SMS oder E-Mail eine Warnung.** «Das System liefert sehr präzise Daten und verfügt über einen leistungsfähigen Akku mit einer Laufzeit von drei bis vier Wochen», sagt Forscher David Atienza.

18 octobre 2011

Le résumé de l'actualité romande

Date: 19.10.2011

LE TEMPS

Santé

Notre cœur sur écoute

L'Ecole polytechnique de Lausanne (EPFL) a annoncé mardi la mise au point d'un appareil de détection des anomalies cardiaques. Les informations seront directement transmises par SMS ou par e-mail au patient et au personnel médical. (ATS)

And tested with SolarImpulse pilots too!

Using ESL-SmartCardia Technology with Solar Impulse

- Monitoring pilots using a SmartPhone as "doctor in the cockpit"

But where did we lose so much possible energy savings? What else can we do?

Smart Wearable Architecture: Power Analysis

- Memory occupies a large percentage of area… And consumes **a lot of power**
  - Power gating not feasible in many wearable applications (retain bio-signal content)

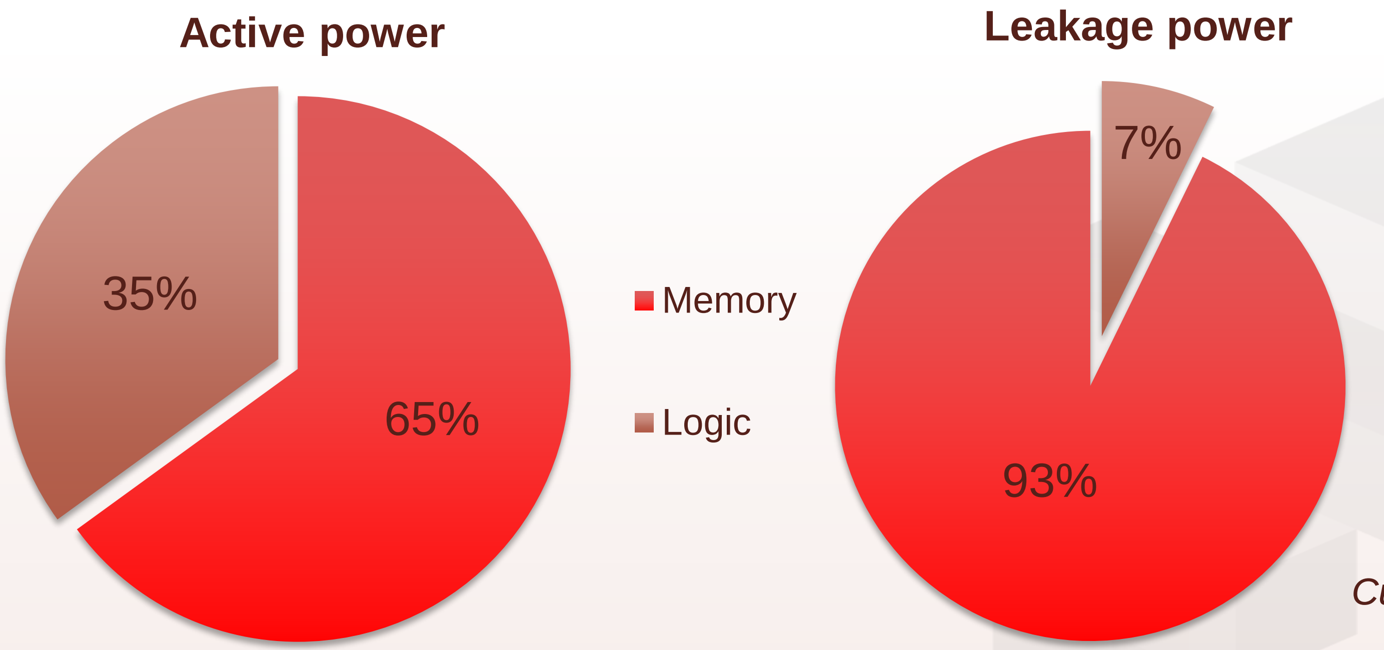

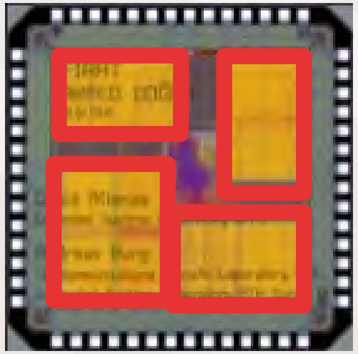

Custom ECG embedded processor

Significant contribution to power consumption through leakage
when logic used at low voltages/frequencies,
reduce memories supply voltage

But Memories Fail at Low VT : Significance-Driven (SD) Computing

- Tech. scaling (starting at 65nm):
  - **Power overhead** for full error coverage: **Power↓ Nr. errors↑**

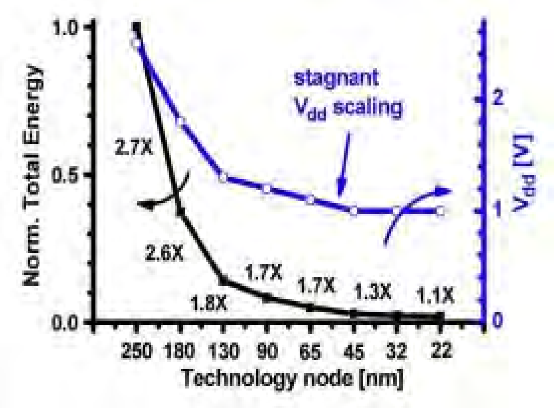

α

- Inherit techniques from video/image processing:
  - Computations do not contribute equally to QoS
  - **Significance-driven**: critical vs non-critical computations

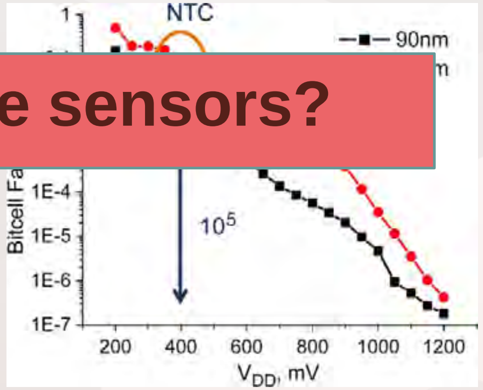

Can SD computing help with memories for wearable sensors?

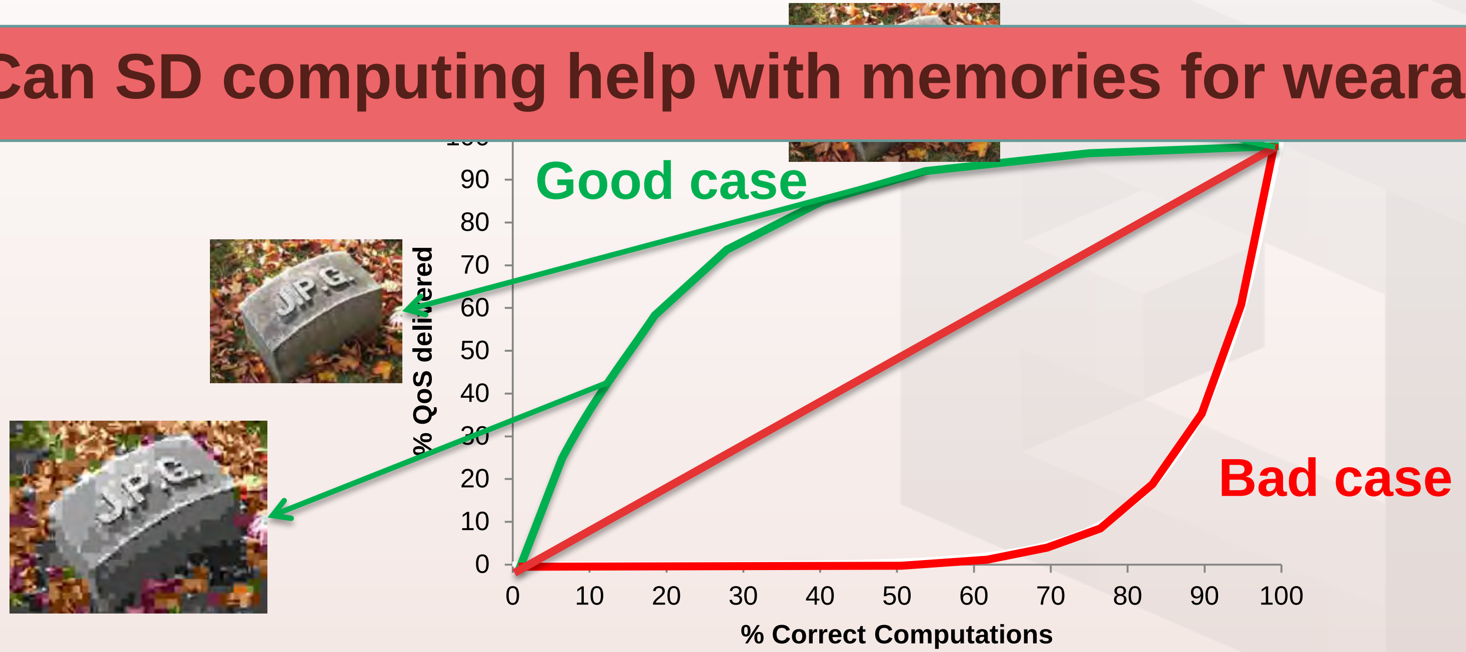

DWT-Based ECG: Selected Significance cases

- Significant data: full protection
  - I: 12.5% significant
  - II: 25% significant
  - III: 37.5% significant
  - IV: 50% significant
- Significance sensitivity analysis
  - Black box approach: output based
  - Inject error (k) and observe faulty output ($\overline{Y}_k$)
- Sensitivity metric: % root-mean square differences:

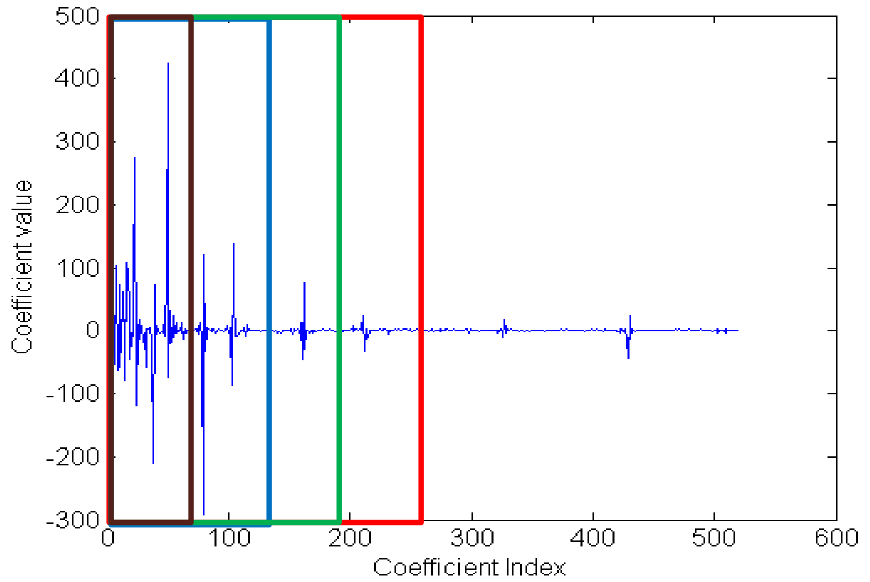

$$PRD(k) = \frac{\left\| Y - \overline{Y}_k \right\|_2}{\left\| Y \right\|_2}$$

Case	PRD
I	63%
II	32%

Only little data percentage is really "significant"

SD computation achieves ~50% lower energy in memories… And doctors monitoring pilots in action did not notice anything! Are we done? Do we have our smart sensor ready as a brain?

- Exploit features of multi-lead ECG **(~2x lifetime)**
  - Specialized instructions for biosignals compression

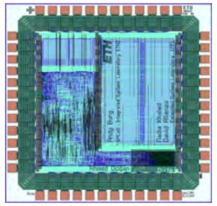

Dicle (umcL 180nm)

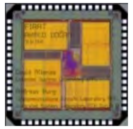

Firat (umcL 90nm)

Hardware: MPSoC fulfils workloads at 50% lower power than single-core wearables, finally smart wearables show true potential!

…PSoC) for biosignals

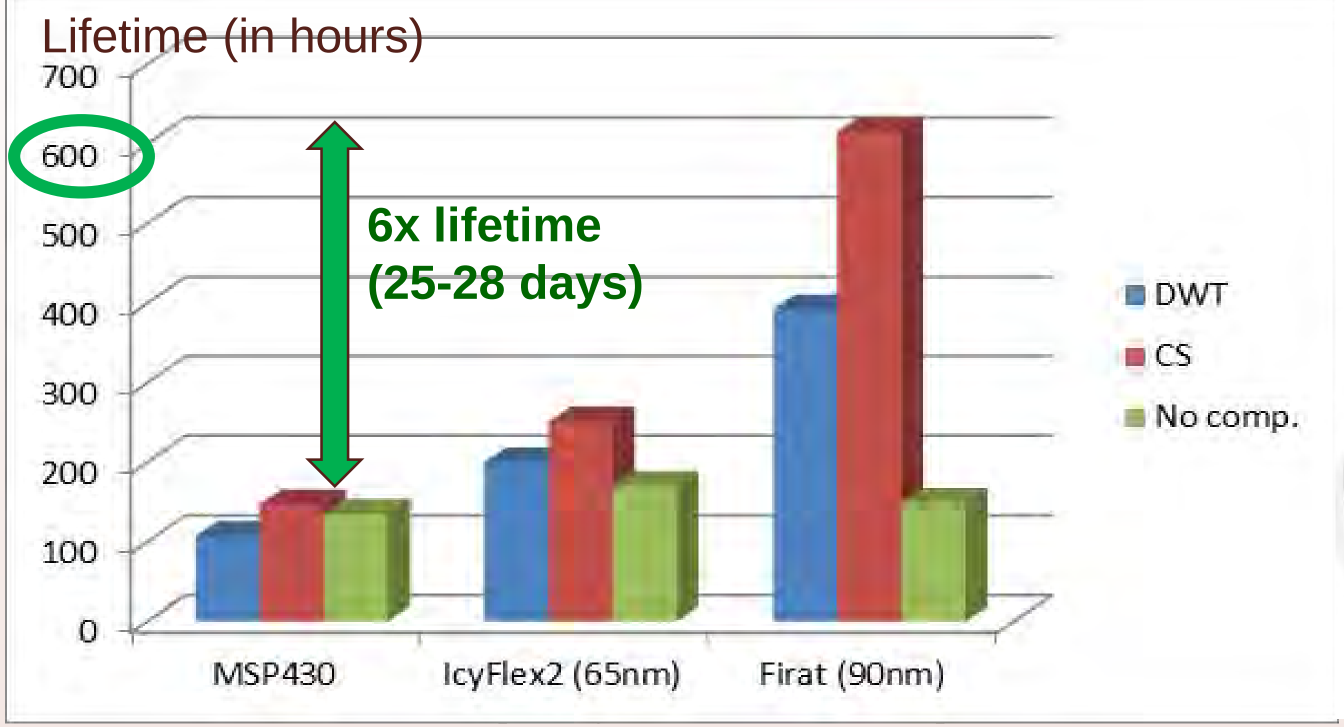

Smart MPSoC Enable New Areas: Multi-Parametric Wearables

- Multiple applications for smart MPSoC wearables, just a few:
  - Accurate sleep apnea tests
  - Epilepsy prediction (non-invasive)
  - Animal health (and stress) monitoring

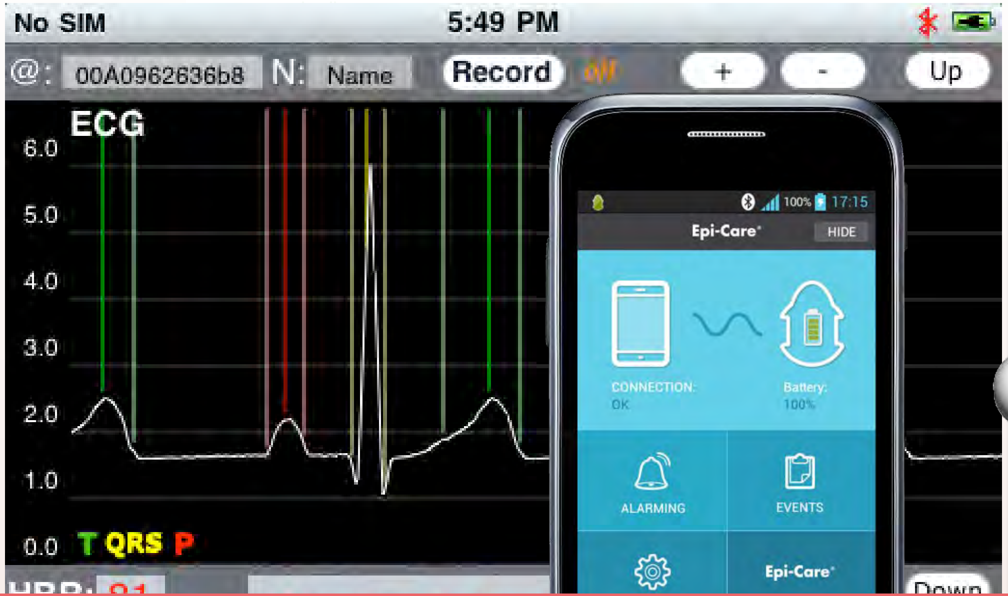

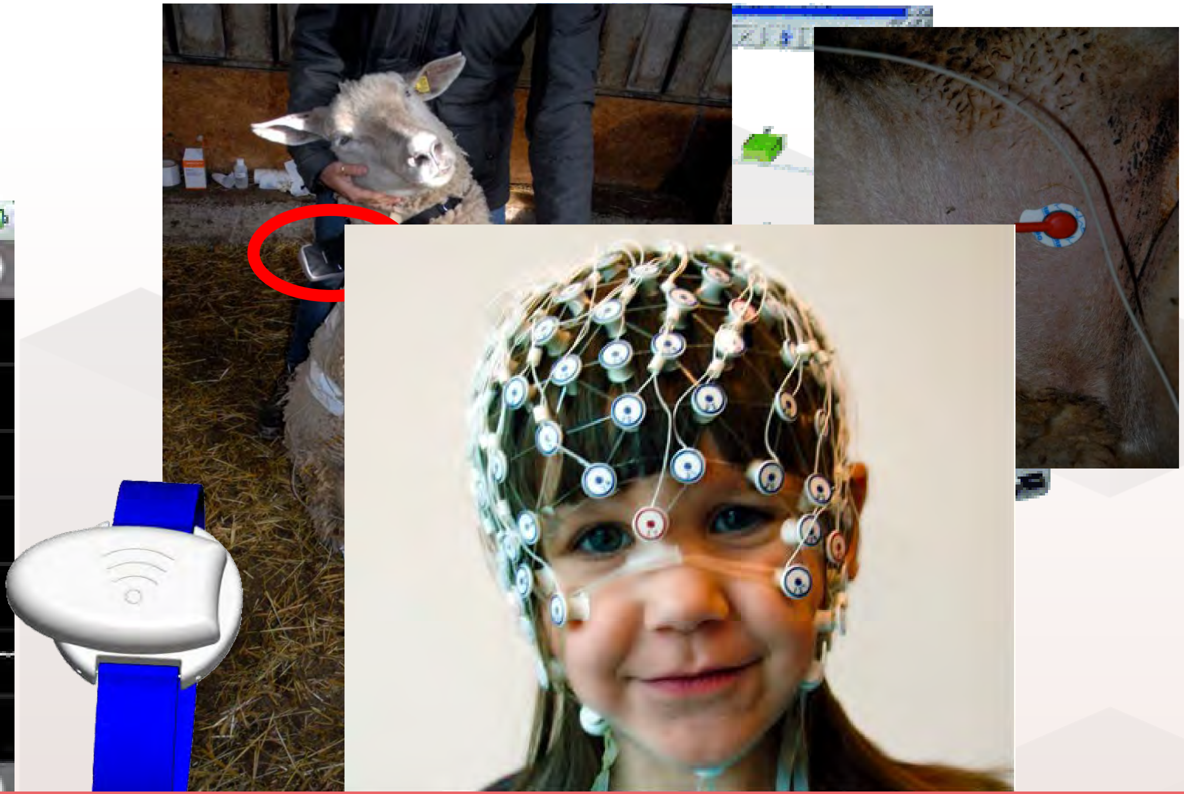

But new dimension: highly variable systems over time,
(1) Adaptability needed, (2) Large volumes of (permanent) data to keep

Revisiting Smart MPSoC Design for Adaptive Wearables

- Current design strategy for smart low-power wearables
  - HW: MPSoC designs + SW: advanced code synchronization
  - Static voltage-frequency scaling **(VFS)** selection

But for adaptive wearables, VFS cannot be exploited anymore, as **more computation is needed**! Need to re-think design strategy

- Power Consumption
  - Power = 1 mW

SOLUTION: Reduce leakage in idle periods

Complete fine-grained power-gating in idle period (but without losing data, as the brain does!)

- Application workload depends on sampling frequency
  - Periodic bursts of computation
  - **Recurrent short idle periods between samples (>90% of time)**

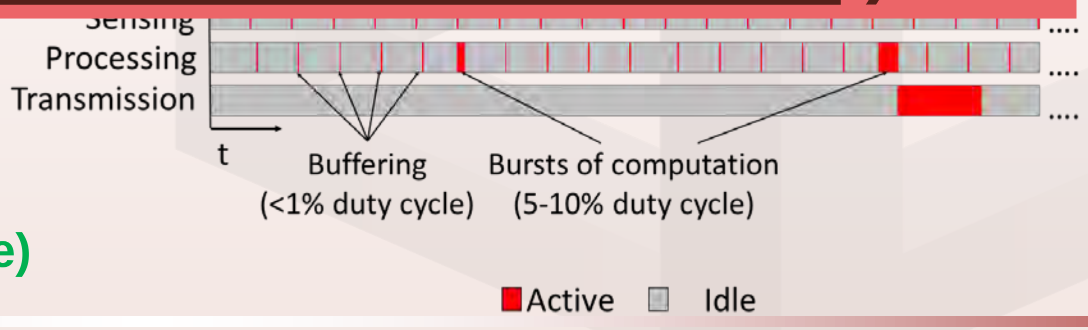

Technological Enablers: Non-Volatile Memory and 3D

1) **Low-voltage Spin-Transfer Torque RAM (STTRAM) NVM**
   - New Memory Subsystem management
     - **Non-Volatile Memory (NVM)**
     - Tiny **volatile page buffers** (I-PB, D-PB) acting as a cache for the NVM
     - **Memory Management Unit (MMU)** + extended **Synchronization Unit**

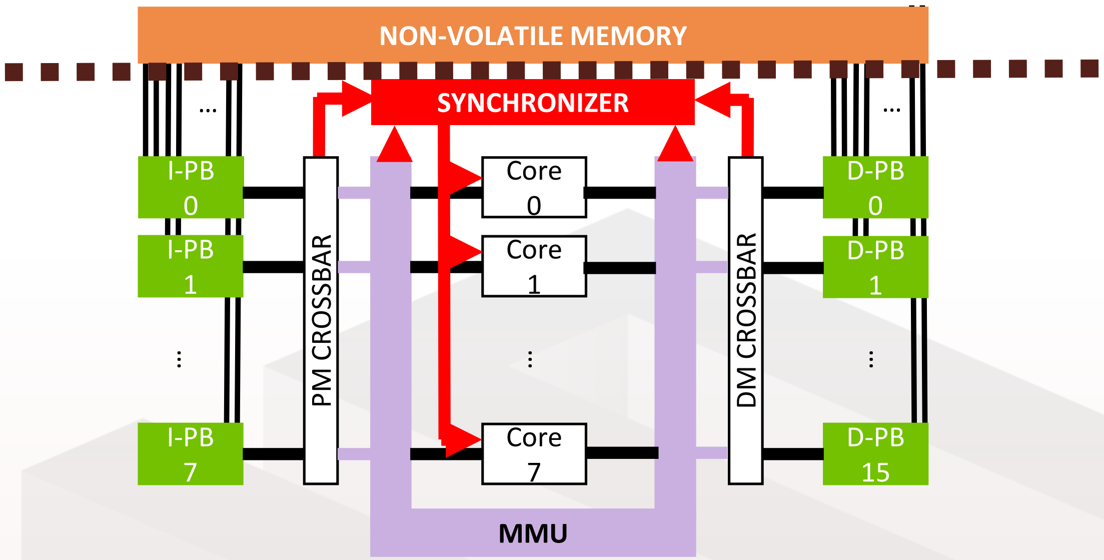

2) **3D monolithic Integration**
   - Multi-tier MPSoC chip
     - Bottom: **Processing Tier**
     - Upper: **NVM Tier**
   - Ultra-dense interconnect: Inter-Level Vias (ILV): NVM ↔ Page buffers
     - 1-cycle page transfer
     - Fast Save/Restore state

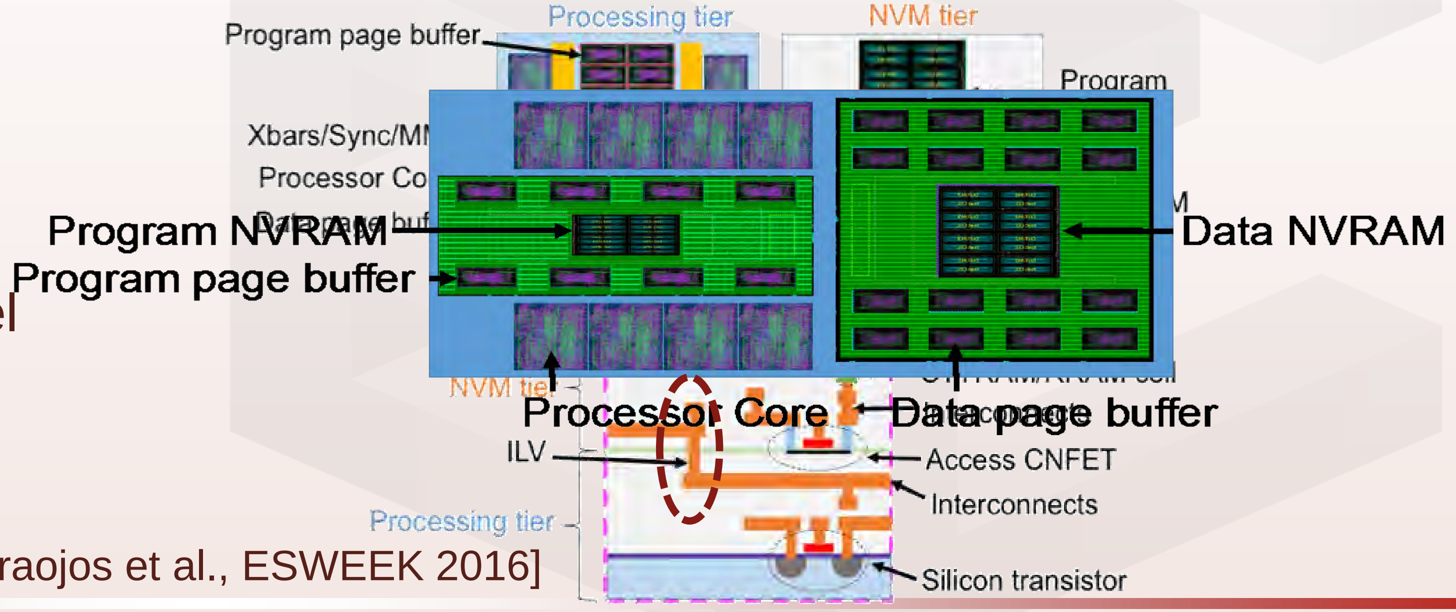

[Braojos et al., ESWEEK 2016]

Power and Area Results: For 3D NVM Smart Wearables

- Power consumption
  - 4 ECG processing apps
  - 2 architectures
    - State of the art **(SOA)**
    - 3D + NVM **(TARGET)**
  - Up to **82% less Leakage**

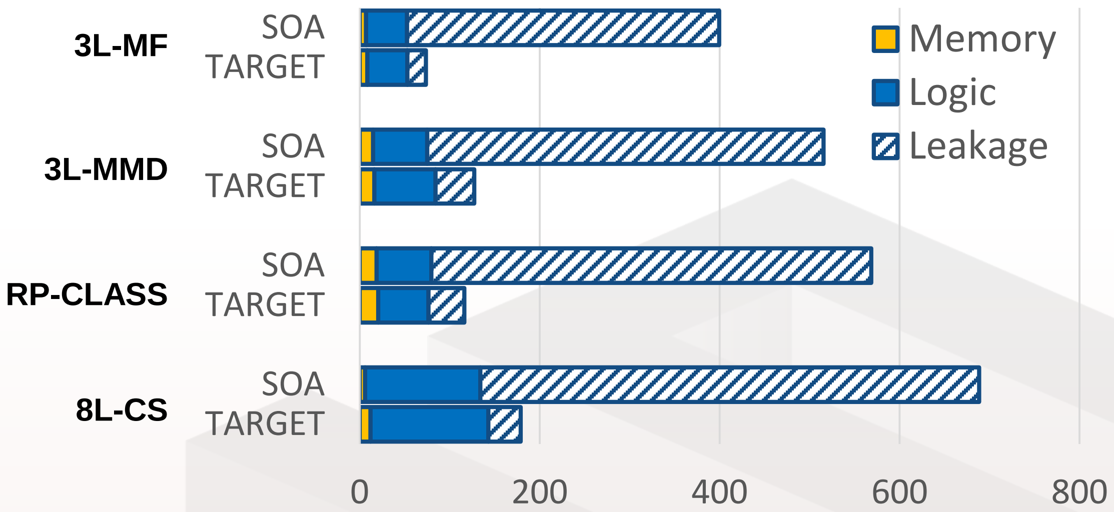

Promising energy (5.42x) and area footprint (5x) reductions with 3D MPSoC NVM architectures for Smart Wearables

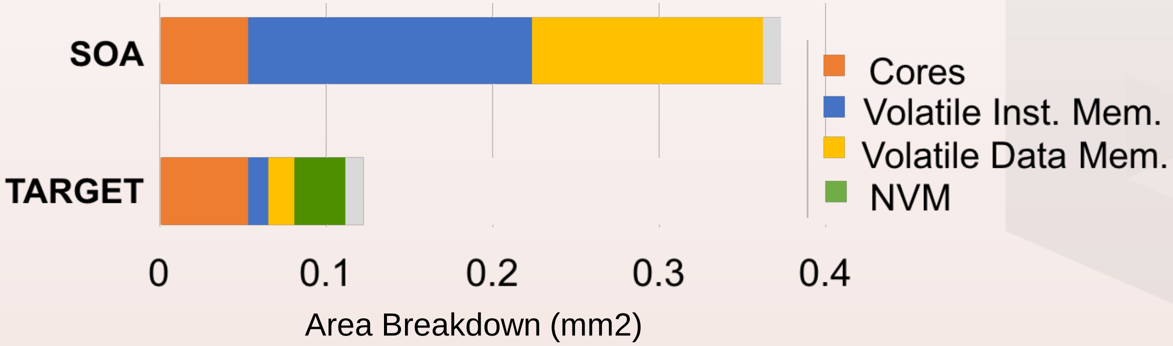

- **Area footprint**
  - Very compact memory subsystem (Inst. and data)
  - **5x less Area Footprint**

[Braojos et al., ESWEEK 2016]

- **Homogeneous MPSoC architecture**
  - Parallel execution
  - Low clock frequency enabled
  - **But not optimized for intensive (repetitive) tasks**
- **Brain training: “HW specialization”**
  - Highly energy efficient
  - Limited configurability (based on iterative training)
  - Application dependent (per domain)
- **Low-power heterogeneous MPSoC reconfigurable architecture**
  - Based on a Coarse-Grained Reconfigurable **Array (CGRA)**
    - High energy efficiency
    - High configurability / flexibility

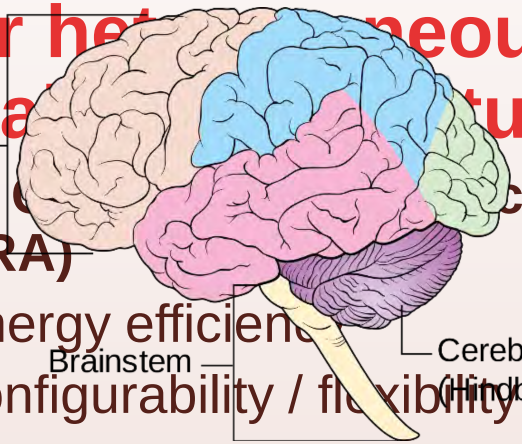

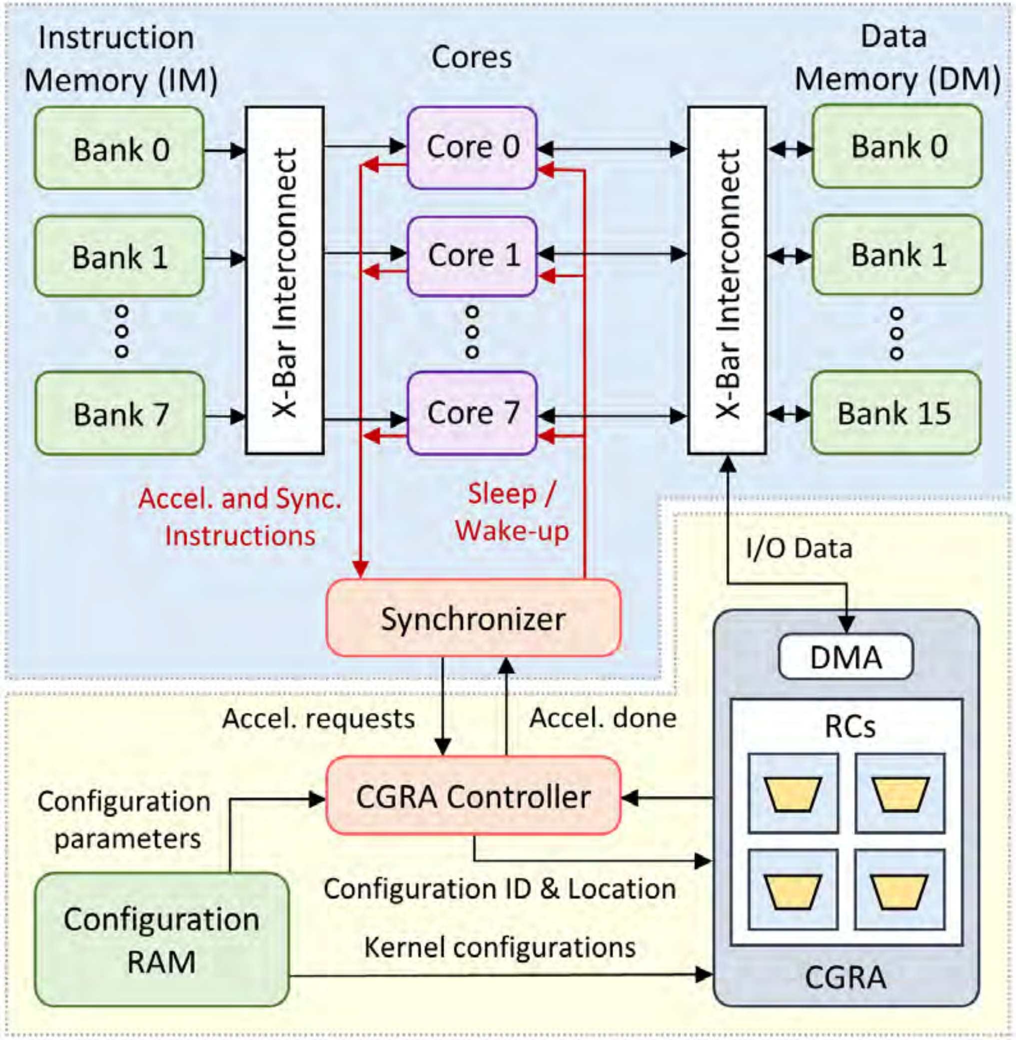

[Duch et al., BioCAS 2016]

Initial Results: Specialized MPSoC is More Energy Efficient

- Complex smart wearable system
  - Filtering (MF+RMS) [Sun et al., 2011],
  - Delineation (MMD) [Rincon et al., 2013]
  - Heartbeat Classifier (RP-CLASS) [Braojos et al., 2015]
- Setup: 5000 Samples – MIT BIH database

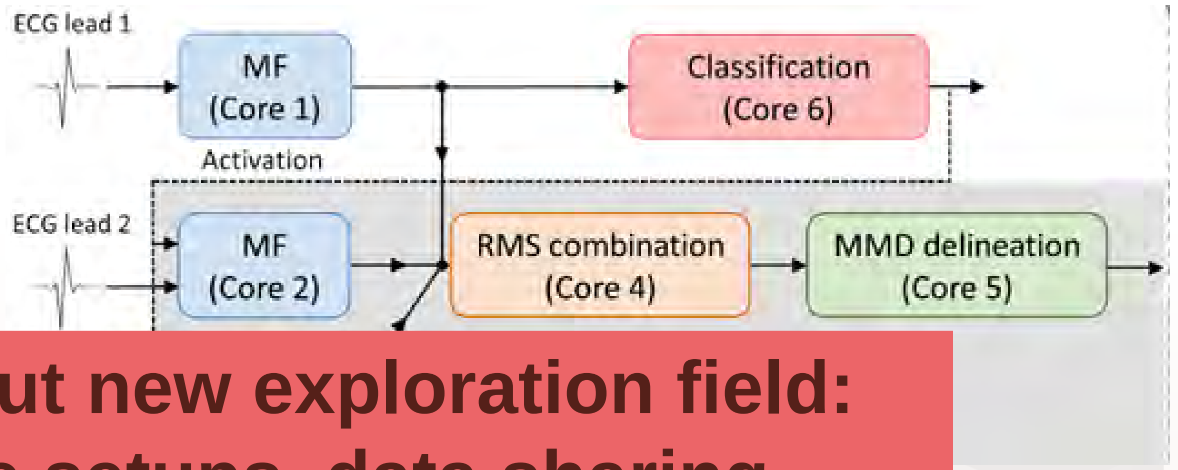

30% energy savings, promising but new exploration field: MPSoC specialization with more setups, data sharing, dynamic reconfiguration…

- considered applications
  - Relevant kernels identifiable in all phases of system
  - CGRA energy compensated by reduction in cores
  - Unused RC cells power-gated

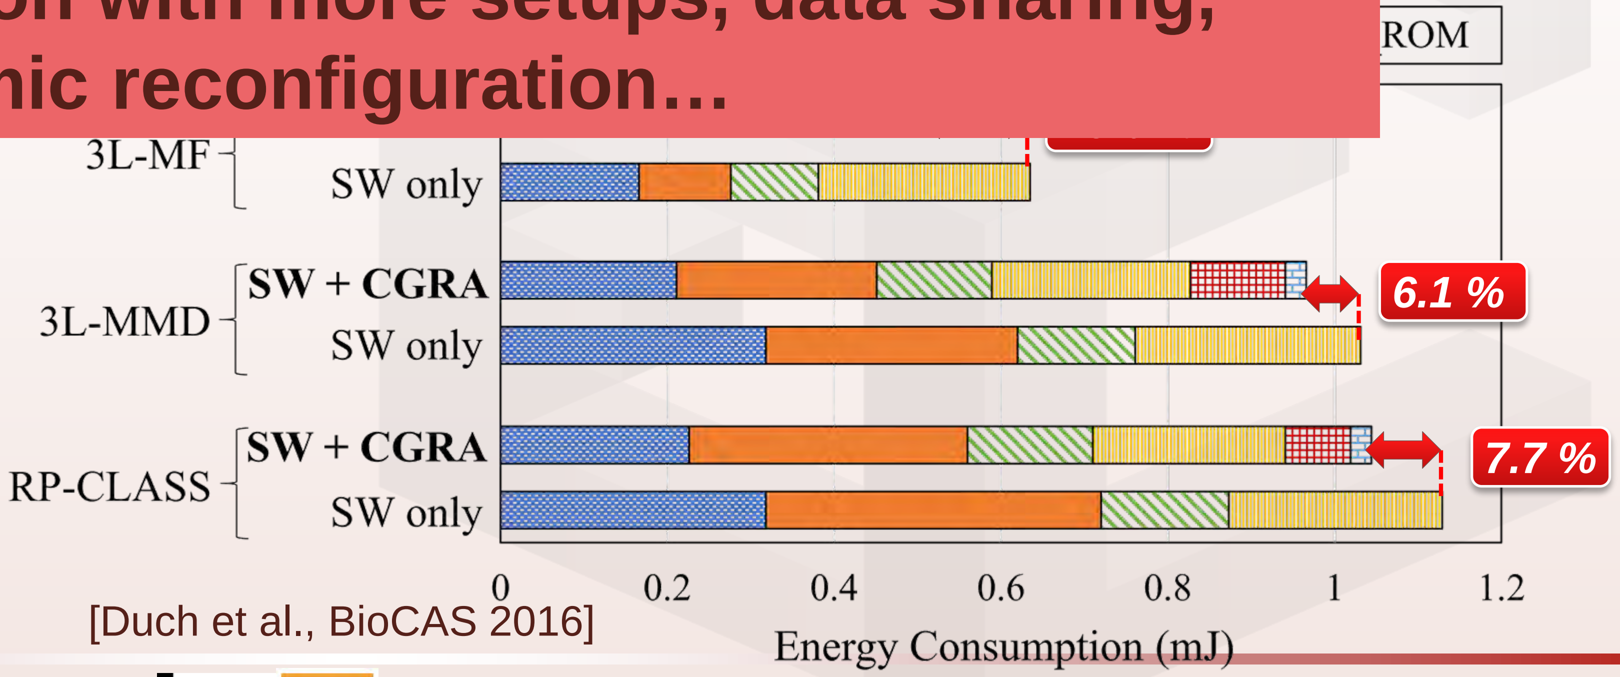

[Duch et al., BioCAS 2016]

- Wearable devices are getting everywhere… "IoT Era" coming (in some domains more than others…)
  - Powerful multi-core architectures and connectivity available
  - But not low power… To be re-designed with care!

- Smart MPSoC wearables needed for IoT
  - Hardware and software equally important
  - Let's get inspiration from biological systems

- Luckily lots of research to reach truly smart wearables, but thanks Mr. Spock for initial idea!
  Tricorder: 1st Smart Wearable! (Sense, compute and record)

david.atienza@epfl.ch

Key References and Bibliography

- ULP WBSN computation optimization and ECG application mapping
  - R. Braojos, H. Mamaghanian, A. Junior, G. Ansaloni, D. Atienza, et al.,"*Ultra-Low Power Design of Wearable Cardiac Monitoring Systems*", Proc. of DAC, 2014.
  - F. Rincon, J. Recas, N. Khaled, D. Atienza, "*Development and Evaluation of Multi-Lead Wavelet-Based ECG Delineation Algorithms for Embedded Wireless Sensor Nodes*", IEEE Trans. on Information Technology in BioMedicine (TITB), Nov. 2011
- Single- vs. multi-core WBSN platform design
  - L. Duch, S. Basu, et al., "*A Multi-Core Reconfigurable Architecture for Ultra-Low Power Bio-Signal Analysis*", Proc. of BioCAS, 2016.
  - R. Braojos, D. Atienza, et al. "*Nano-Engineered Architectures for Ultra-Low Power Wireless Body Sensor Nodes*", Proc. of CODES-ISSS, 2016.
  - R. Braojos, I. Beretta, G. Ansaloni, D. Atienza, "*Hardware/Software Approach for Code Synchronization in Low-Power Multi-Core Sensor Nodes*", Proc. of DATE, 2014.
  - A. Y. Dogan, J. Constantin, M. Ruggiero, D. Atienza, et al., "*Multi-Core Architecture Design for Ultra-Low-Power Wearable Health Monitoring Systems*", Proc. DATE, 2012.
- CS-based ECG delineation and implementation
  - H. Mamaghanian, N. Khaled, D. Atienza, P. Vandergheynst, "*Compressed Sensing for Real-Time Energy-Efficient ECG Compression on Wireless Body Sensor Nodes*", IEEE Trans. on Biomedical Engineering (TBME), 2011
  - K. Kanoun, H. Mamaghanian, N. Khaled, D. Atienza, "*A Real-Time Compressed Sensing-Based Personal Electrocardiogram Monitoring System*", Proc. DATE, 2011.

Key References and Bibliography

- ULP biosignal analysis and optimization
  - R.Braojos, I. Beretta, G. Ansaloni, D. Atienza, "*Early Classification of Pathological Heartbeats on Wireless Body Sensor Nodes*", MDPI Sensor, Dec. 2013.
  - R. Braojos, G. Ansaloni, D. Atienza, "*A Methodology for Embedded Classification of ECG Beats Using Random Projections*", Proc. of DATE, 2013.
  - H. Mamaghanian, N. Khaled, D. Atienza, P. Vandergheynst, "*Design and Exploration of Low-Power Analog to Information Conversion Based on Compressed Sensing*", IEEE Journal on Emerging and Selected Topics in Circuits and Systems (JETCAS), Sept. 12.
  - N. Boichat, N. Khaled, F. Rincon, D. Atienza, "*Wavelet-Based ECG Delineation on a Wearable Embedded Sensor Platform*", Proc. BSN, 2009.
- Significance-Driven Computing on WBSN
  - M. Sabry, D. Atienza, F. Catthoor, "*OCEAN: An Optimized HW/SW Reliability Mitigation Approach for Scratchpad Memories in Real-Time SoCs*", ACM TECS, Apr. 2014
  - G. Karakonstantis, M. Sabry, D. Atienza, A. Burg, "*A Quality-Scalable Spectral Analysis System for Energy Efficient Health Monitoring*", Proc. of DATE, 2014.
  - M. Sabry, G. Karakonstantis, D. Atienza, A. Burg, "*Design of energy efficient and dependable health monitoring systems under unreliable nanometer technologies*", Proc. of BodyNets, 2012.