Beyond the System-on-a-Chip

From Integrated Platforms to Distributed Systems

GRC ETAB Summer Study – June 2009

Jan M. Rabaey, Donald O. Pederson Distinguished Professor
Director Gigascale Systems Research Center (GSRC)
Scientific Co-Director Berkeley Wireless Research Center (BWRC)
University of California at Berkeley

Dominating Themes in Integrated Electronic System Design over the Past Decade(s)

- Increased Complexity (System-on-a-Chip)
- Concurrency
- Energy
- Reliability/Resiliency

While major progress has been made, many challenges still remain, especially in light of:

- Reduced benefits of technology scaling
- Effects of nanometer scale devices

The GSRC System-Design Roadmap

Addressing major challenges in embedded system design in the next two decades, as formulated in the GSRC System-Design Roadmap

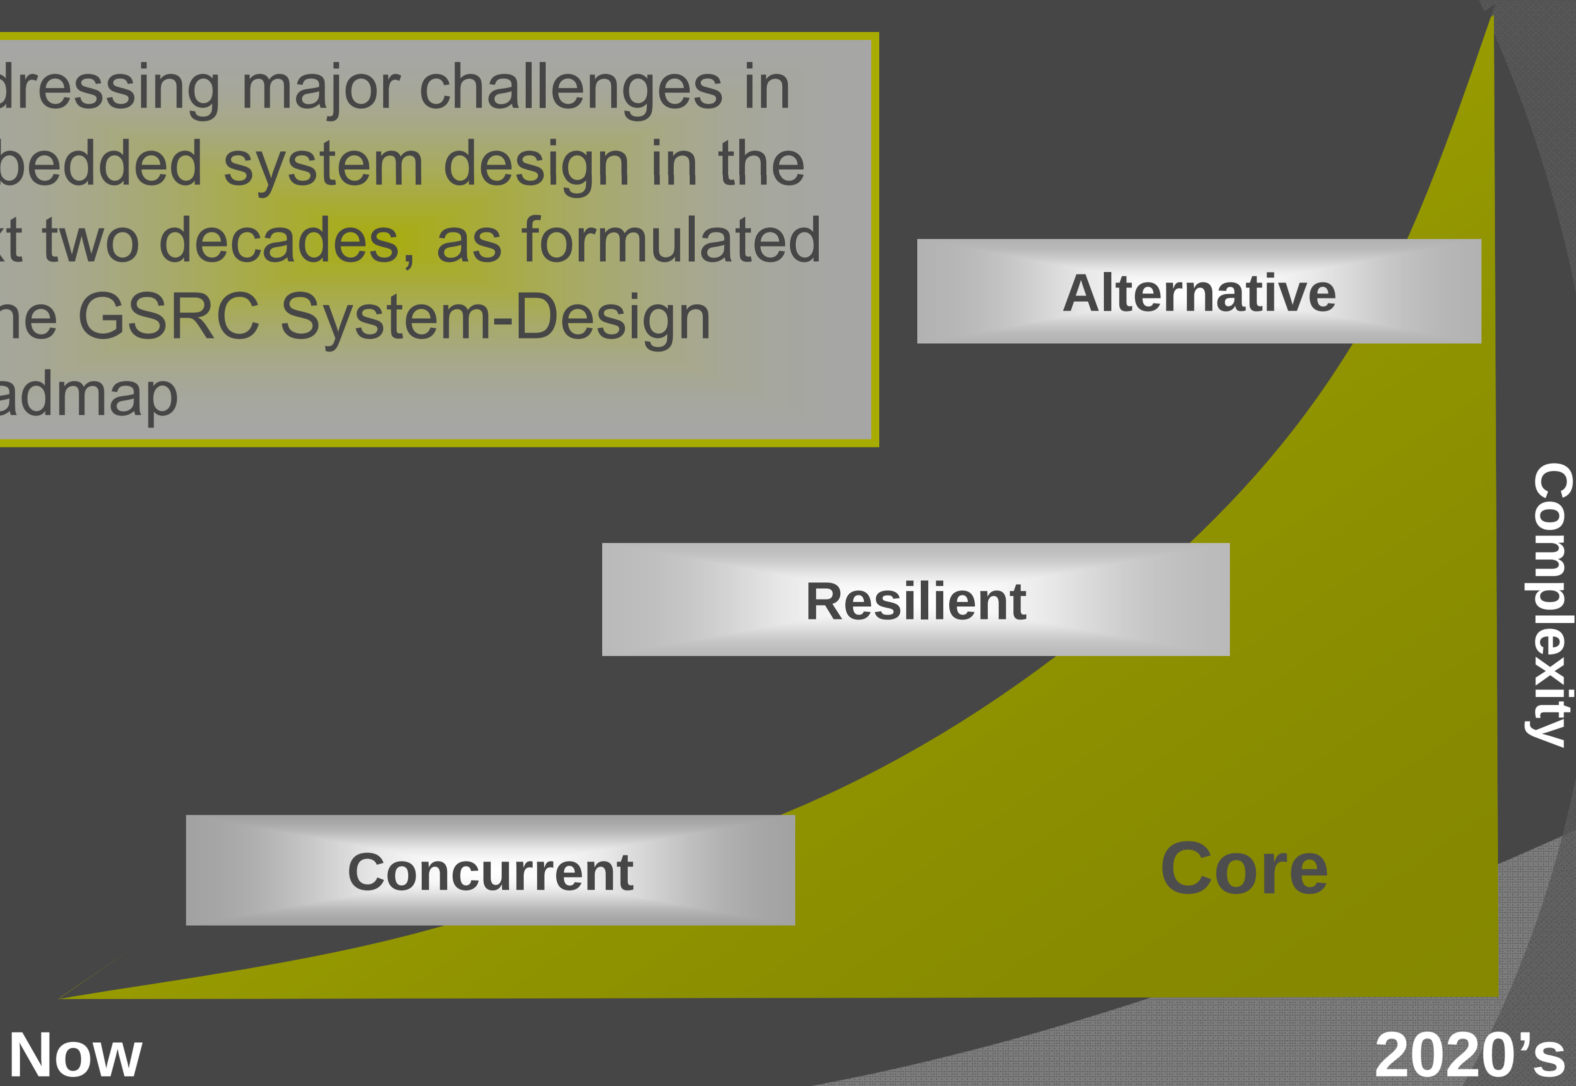

Themes for the Next Decade

- Beyond the system-on-a-chip
  - It's a connected world
- The problems have not changed – just the scales
  - Energy
  - Concurrency
  - Reliability/resiliency
  - Complexity

It's A Connected World

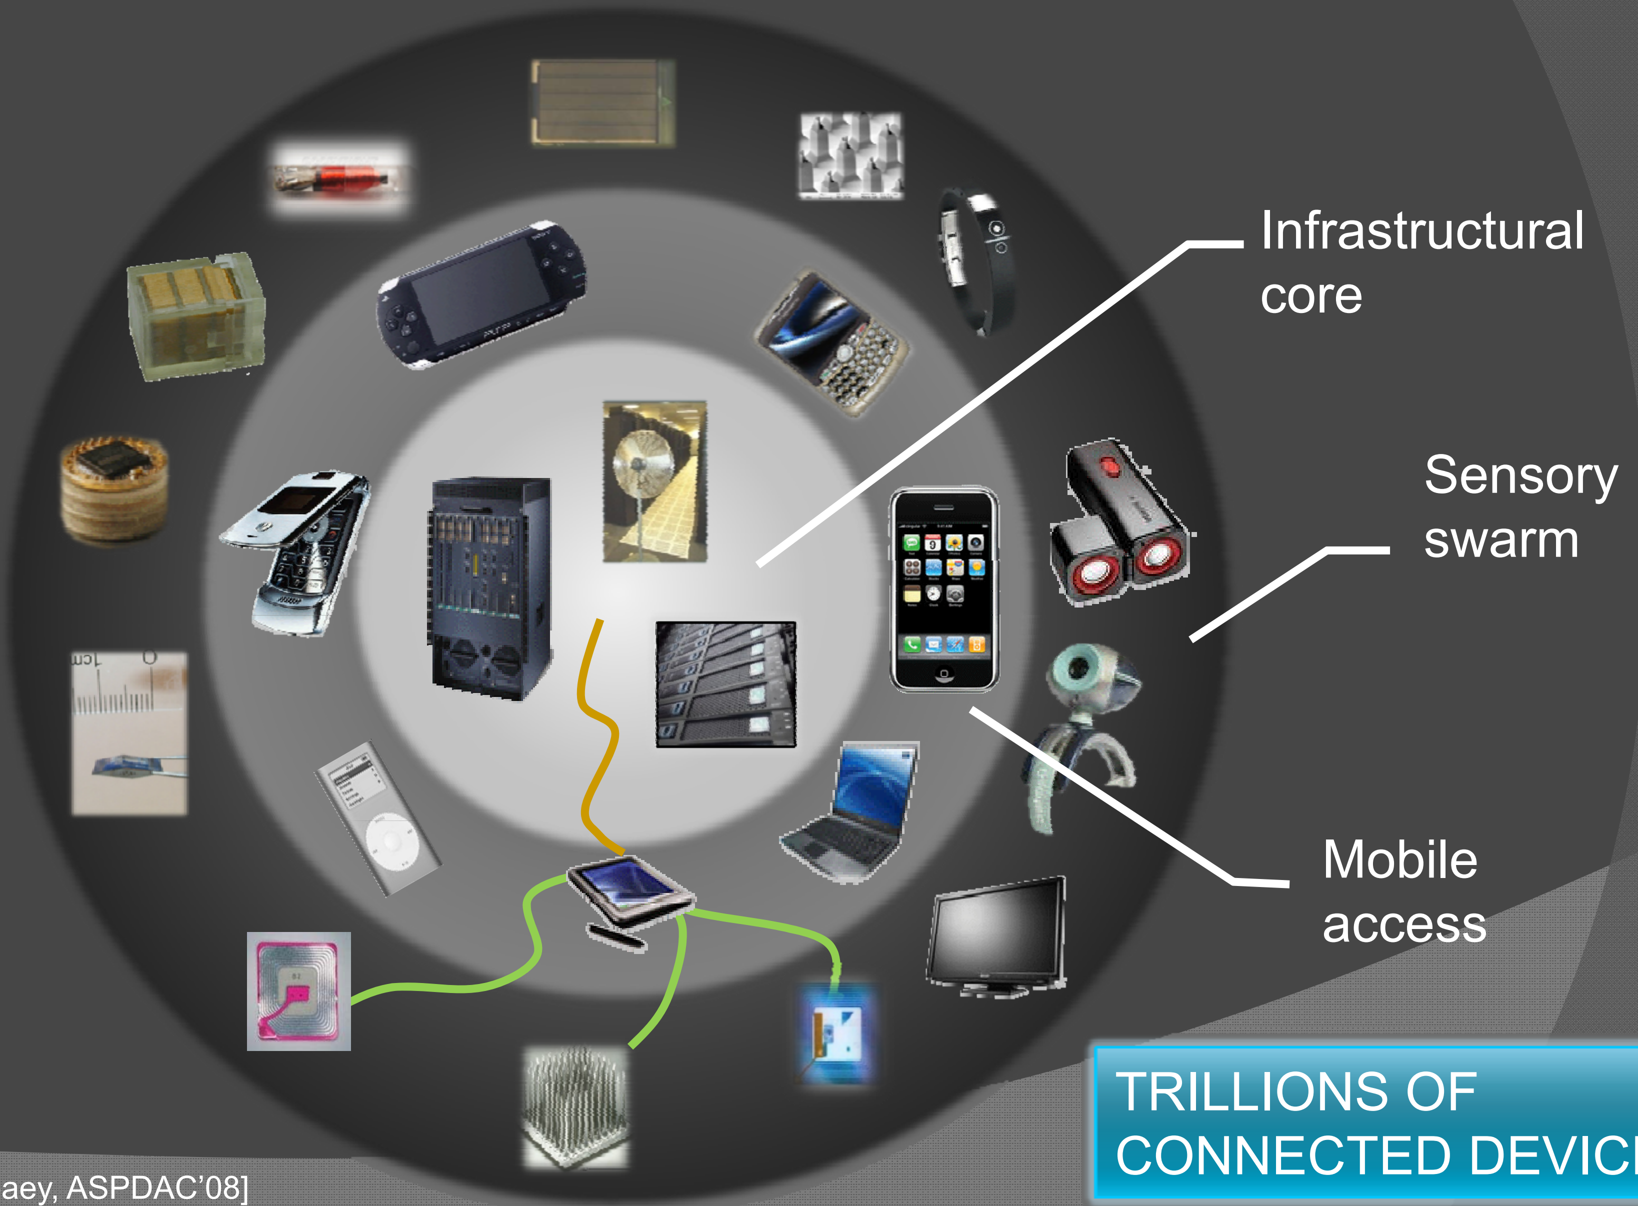

TRILLIONS OF CONNECTED DEVICES

[J. Rabaey, ASPDAC'08]

The Birth of *Societal IT Systems**:

Looking Beyond the Devices



Complex collections of sensors, controllers, compute and storage nodes, and actuators that work together to improve our daily lives




*Also known as SiS

The Technology Gradients

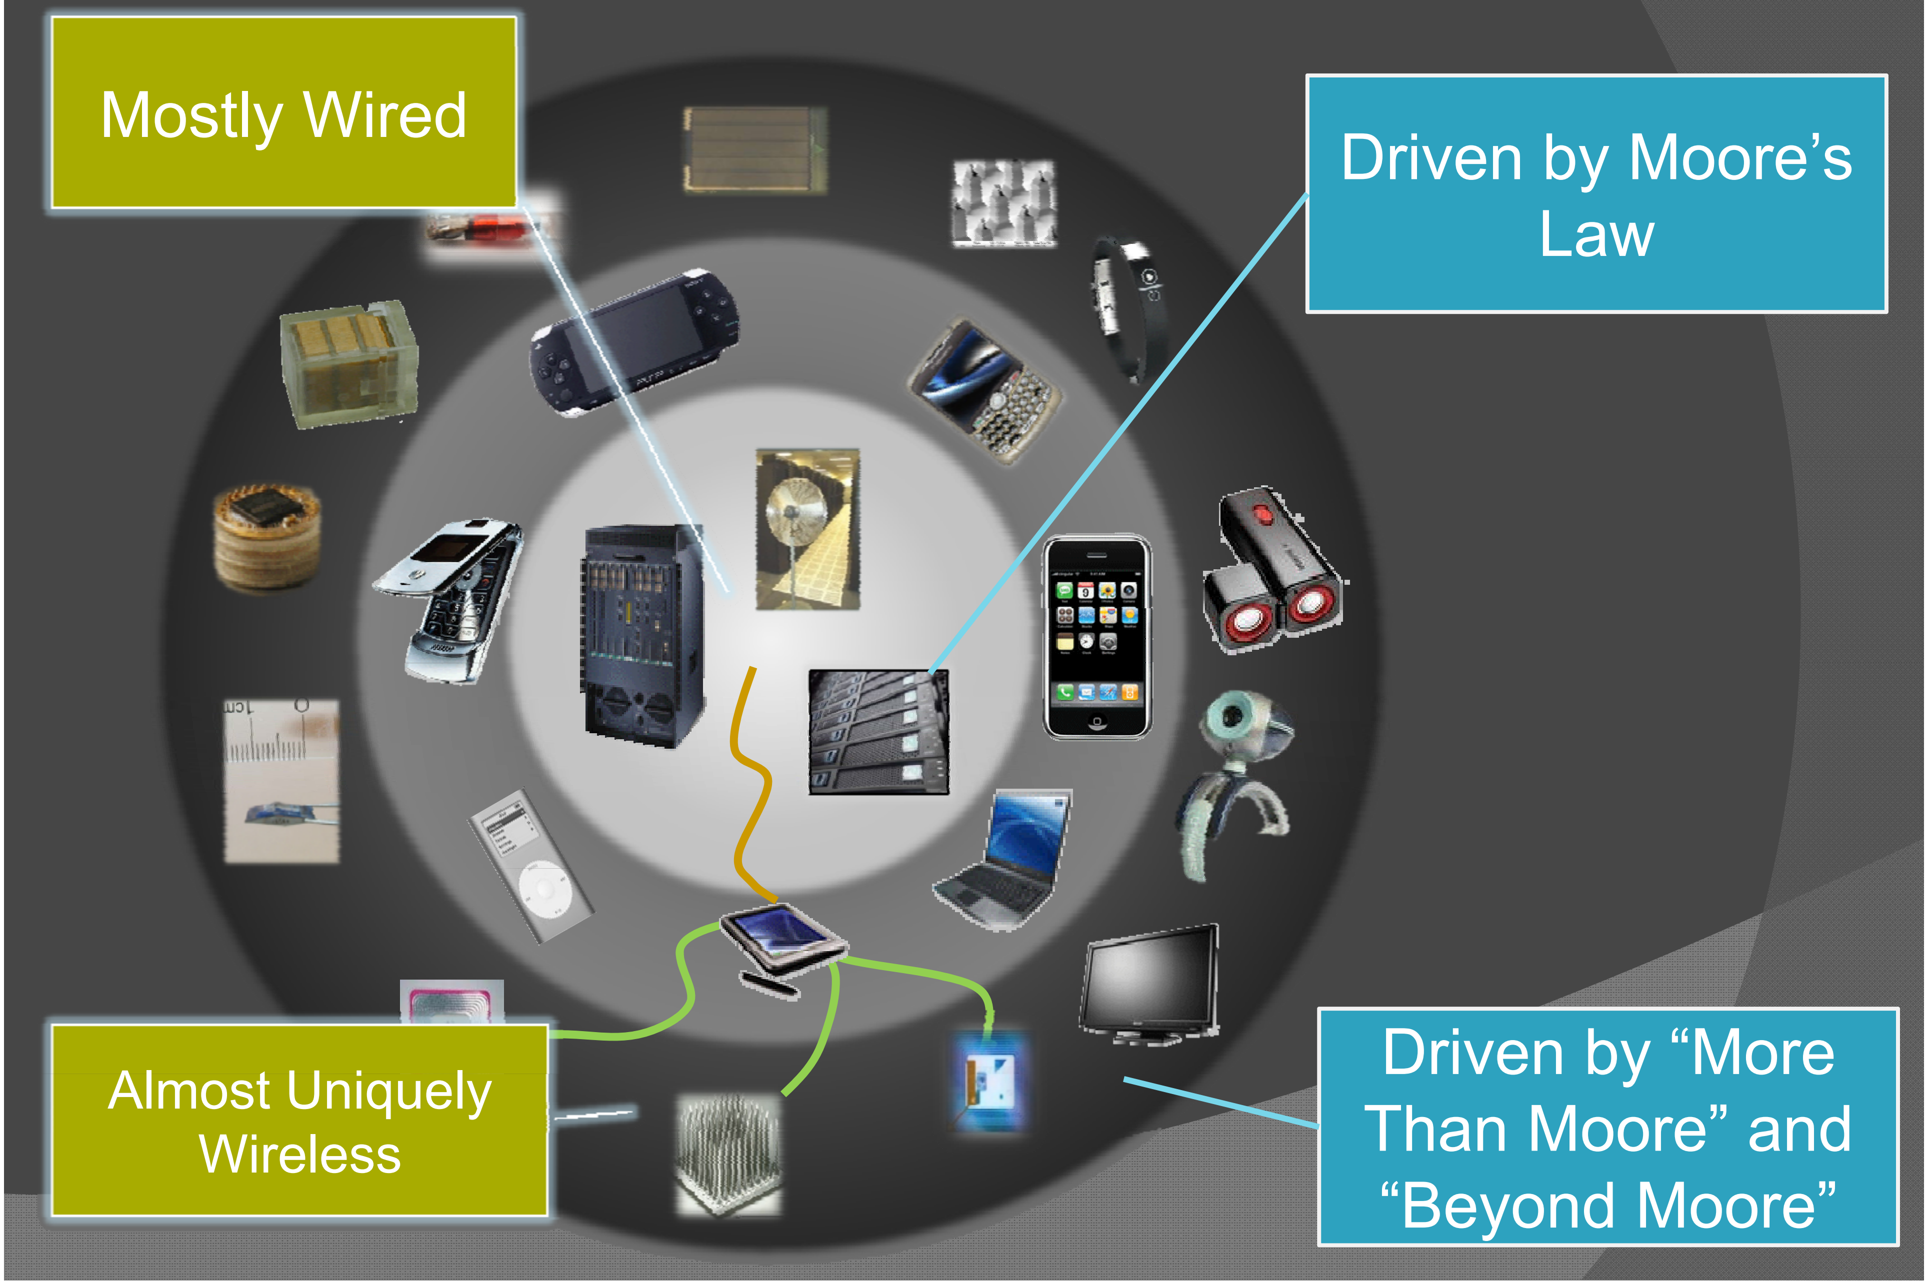

What Does This Mean For Design?

The n-furcation of the design and technology space

- The infrastructuctional core
- The mobile
- The sensory swarm
- Overlaying it all: distributed applications

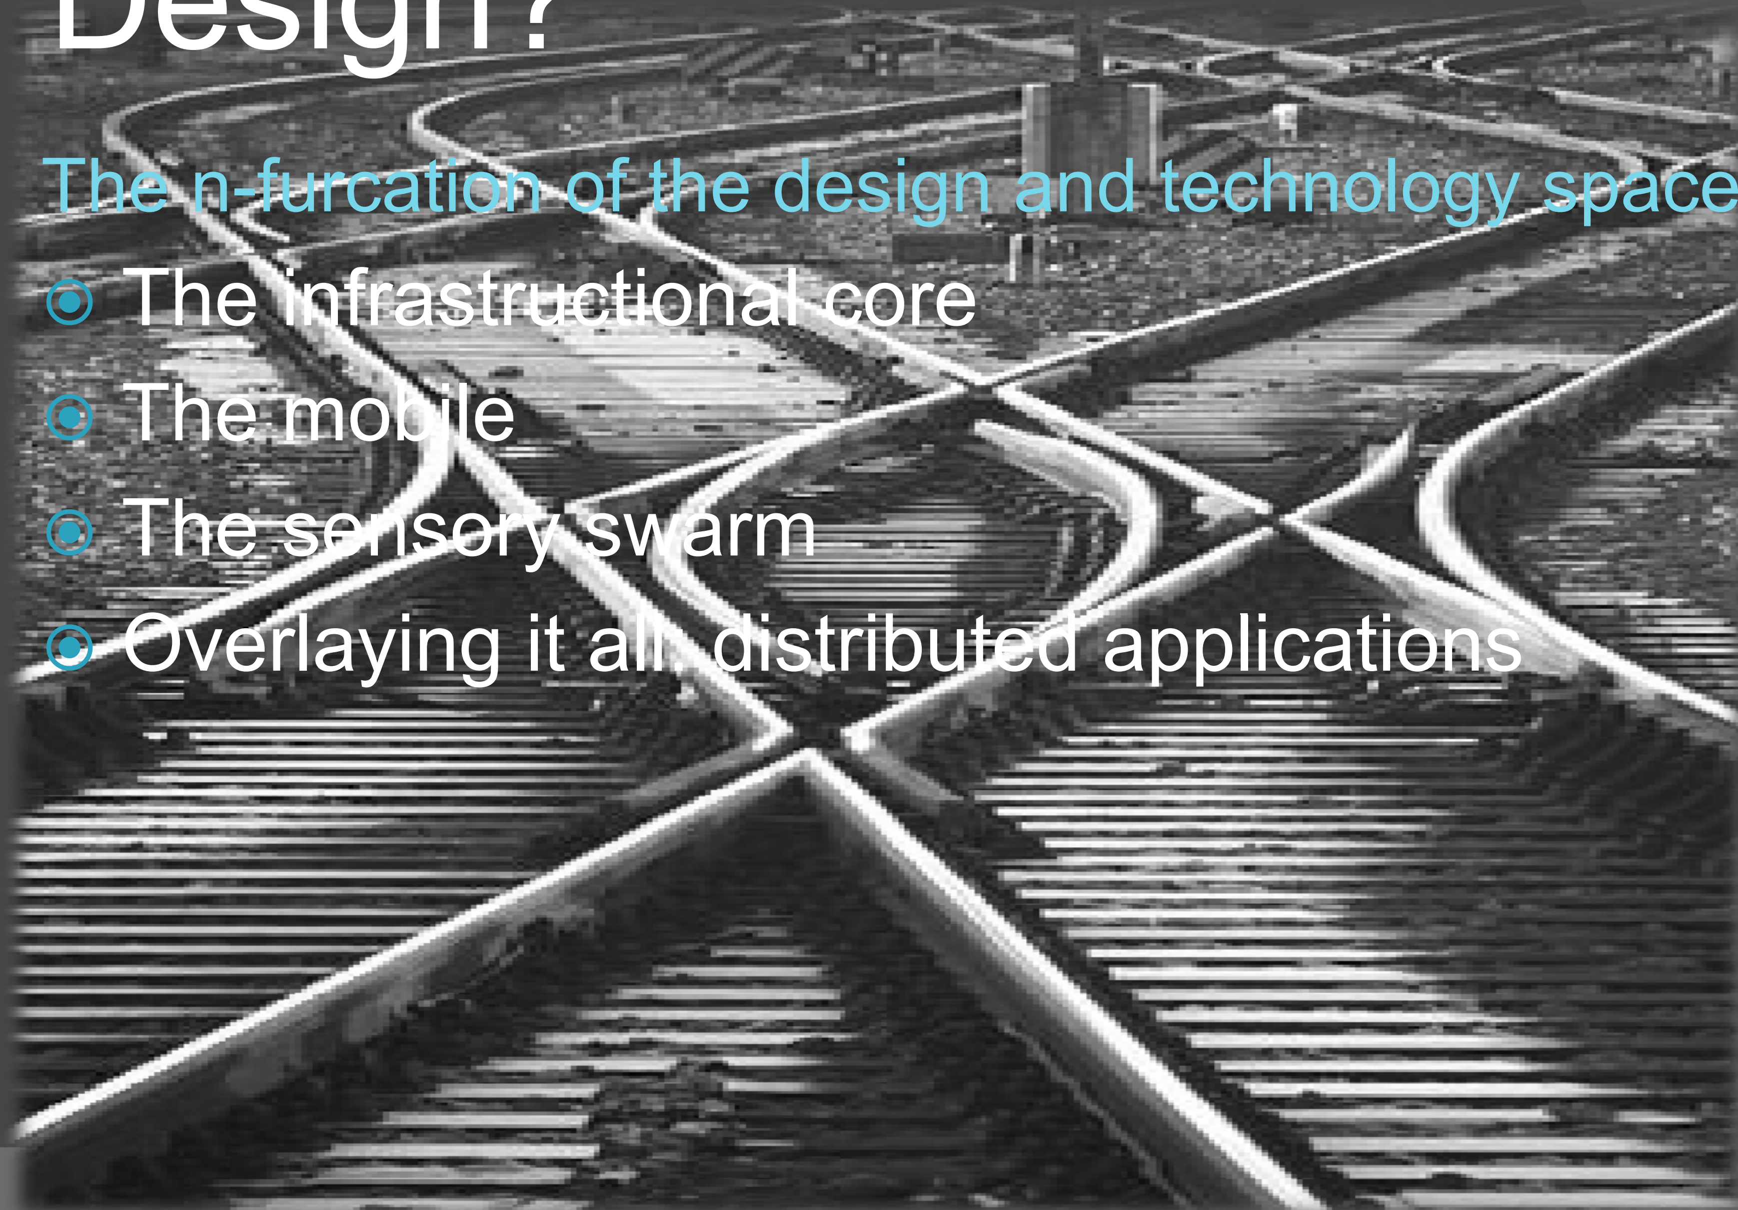

It Is All About Energy …

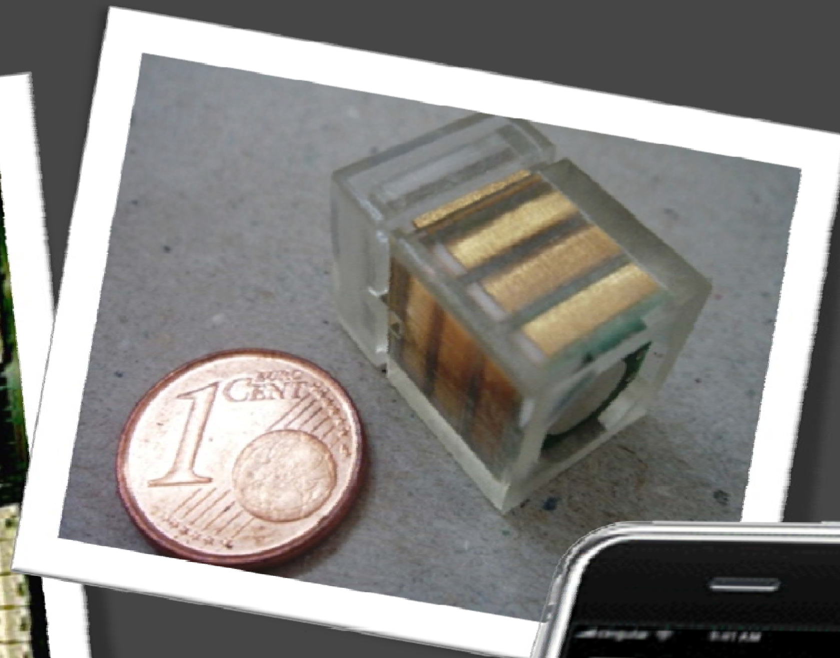

The Infrastructural Core

- An unending appetite for computing, storage and fast networking
- Performance is key, yet power of equal importance

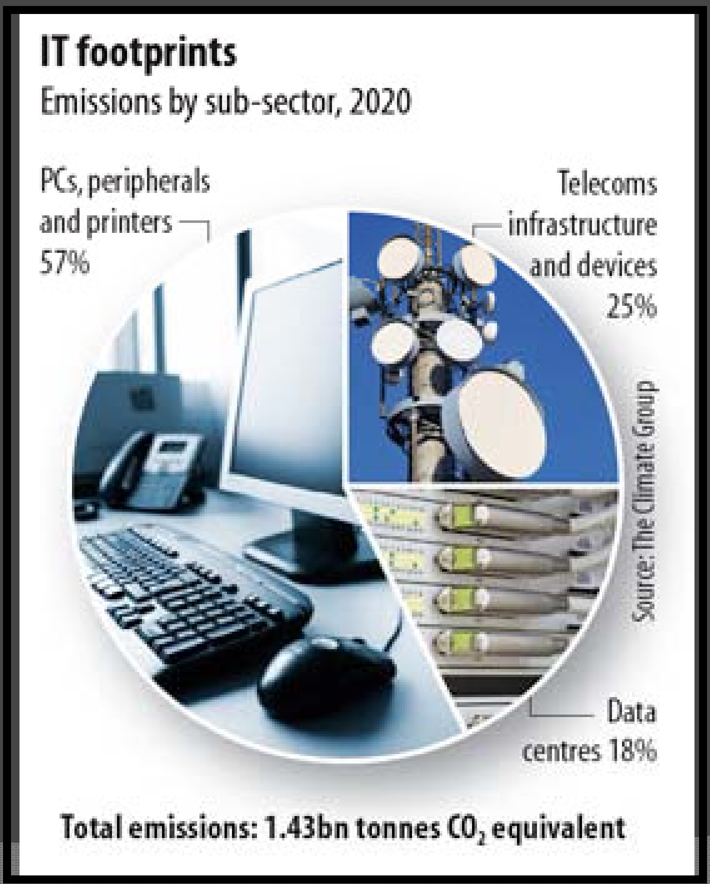

2007 Worldwide IT carbon footprint:
2% = 830 m tons CO_2
Comparable to the global aviation industry

Expected to grow to 4% by 2020
Data centers the fastest growing component

Large Scale Computing and Power

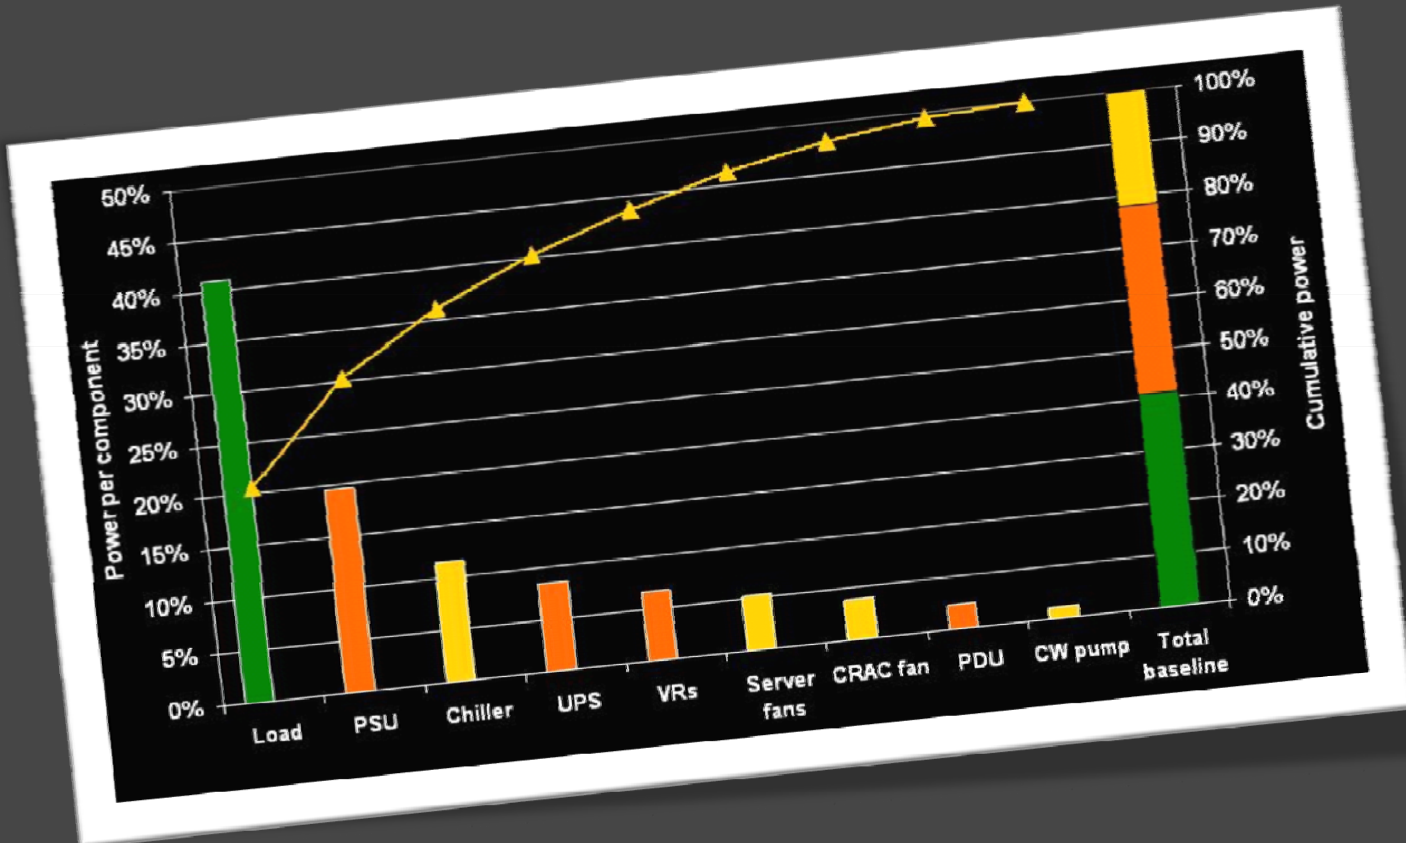

Power budget distributed over many components:
Computing, networking, storage
(as well as power provision and cooling)

The Opportunity:
"Energy-Proportional Computing"

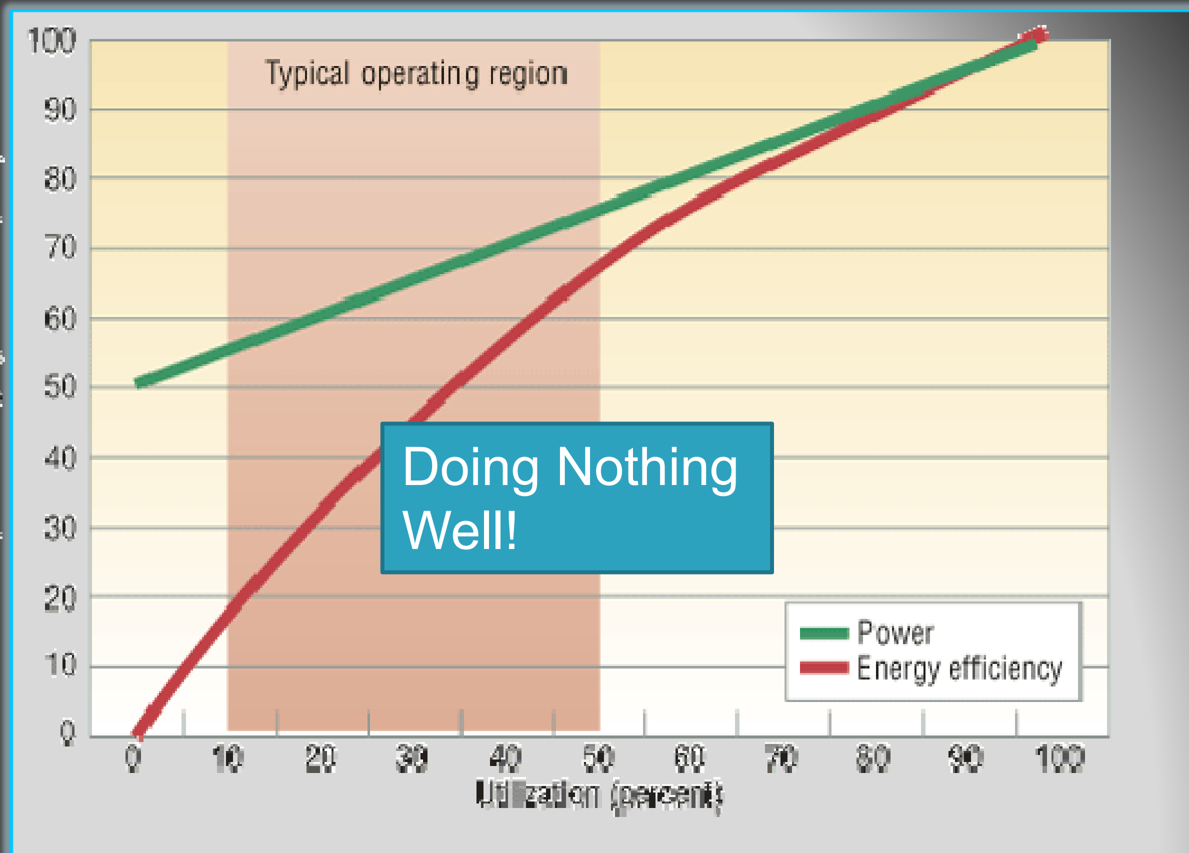

While improving efficiency is essential, reducing waste has larger impact in short term

[Barroso, Holzle, 2007]

Major Opportunity is in Power Management

Requires Top-Down System Level Solution Addressing all System Components

The Mobile Access Device

- The cell phone and its descendants as the “personal” communication and computation device of choice
- Bringing together many different functionalities
- Becoming a “base station” in itself

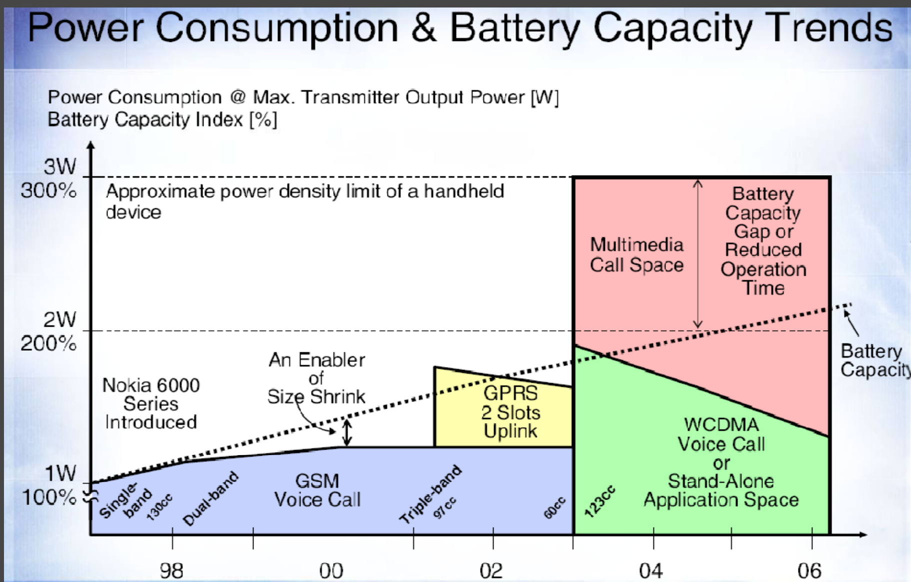

Increasing performance demands under fixed power budget

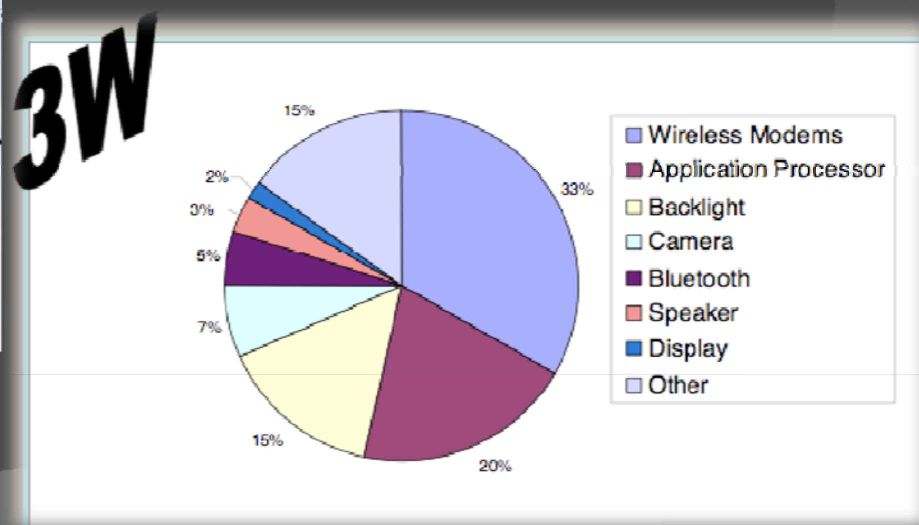

Y. Nuevo, ISSCC 04

Increasing Mobile Energy Efficiency

- Multi-core platforms only partial answer
  - Energy efficiency quickly saturates
- Some improvement with architectural innovations
  - Heterogeneity, accelerators, SoC,…
- The real opportunities
  - A system perspective
  - The mobile as the home of the user interface

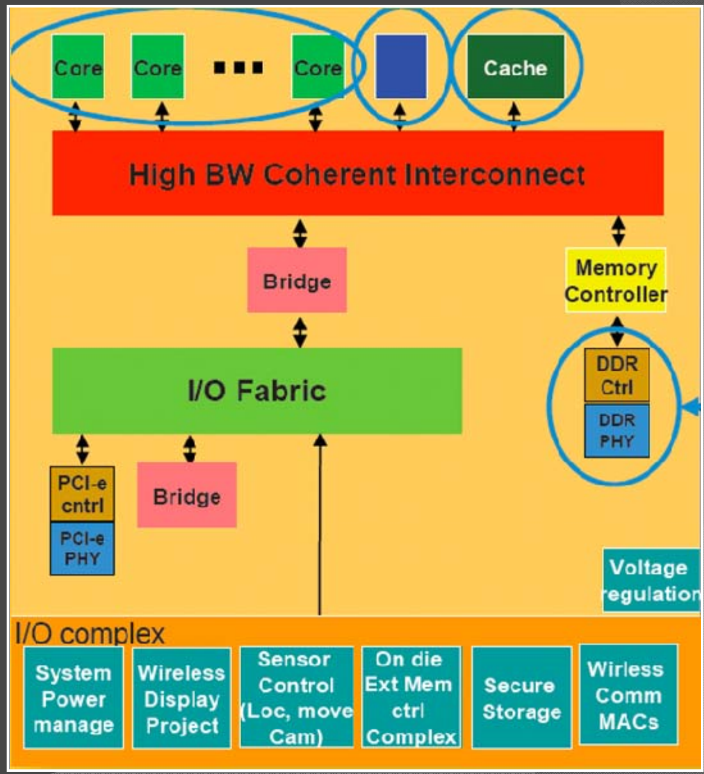

Mobile μProc Anno 2015
(Courtesy A. Peleg, Intel)

The Mobile as the Home of the User Interface

- Innovative interaction paradigms between user and information are gaining ground (e.g. WII)
  - Recognition, Mining, Synthesis (RMS)
- Opens the door for innovative energy-efficient algorithms and architectures
  - Often it is ok to make errors!
  - E.g. ERSA Architecture resilient up to 10^{16} FITS

ERSA Architecture
(Mitra – Stanford)

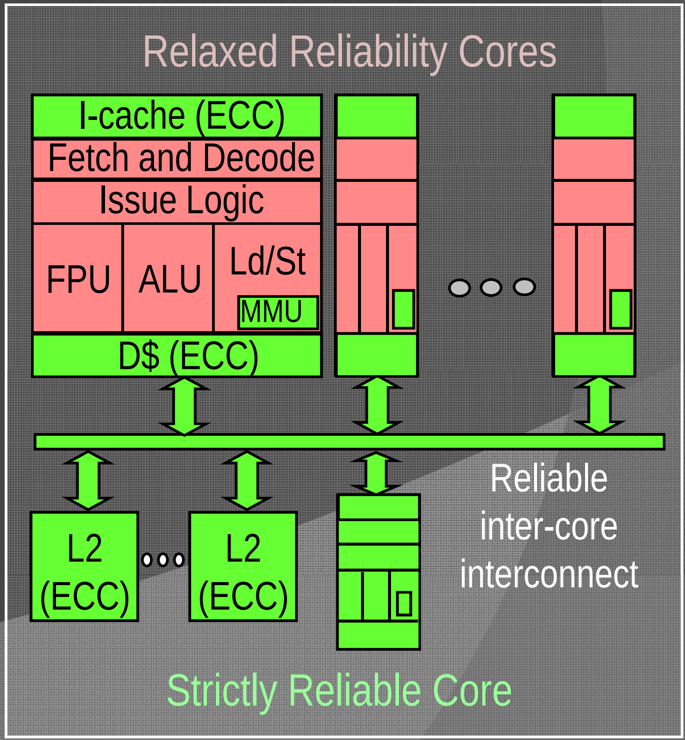

The Sensory Swarm

"Adding senses to the Internet"

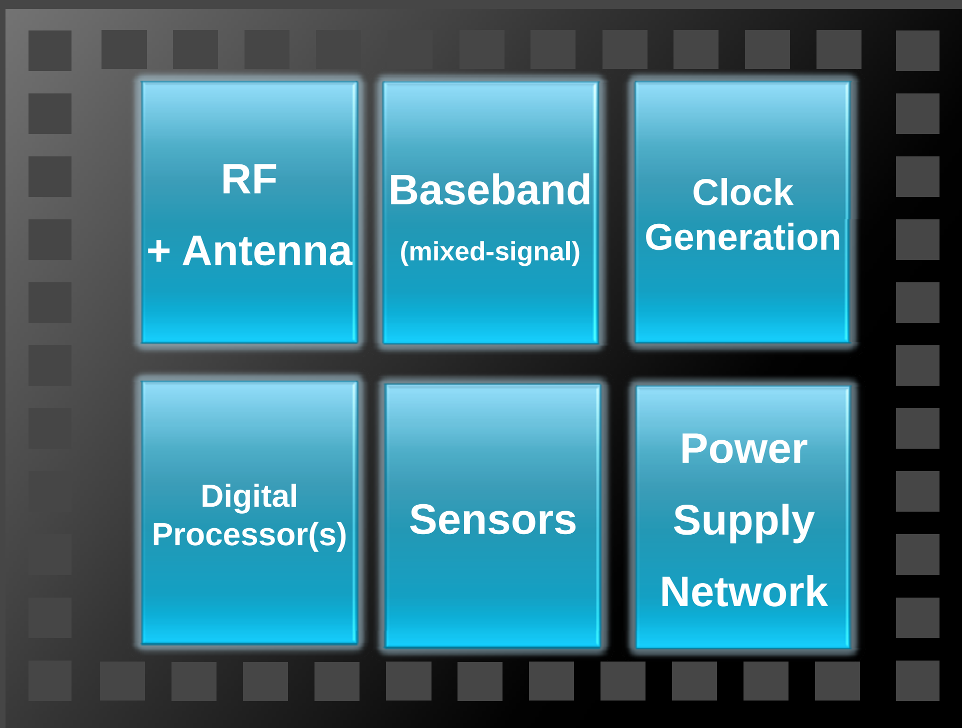

"Disappearing electronics"

- Low-cost
- Miniature size
- Self-contained from energy perspective

Major Progress Over Past Years

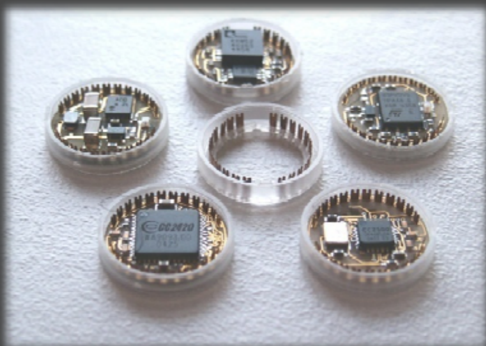

Philips Sand module

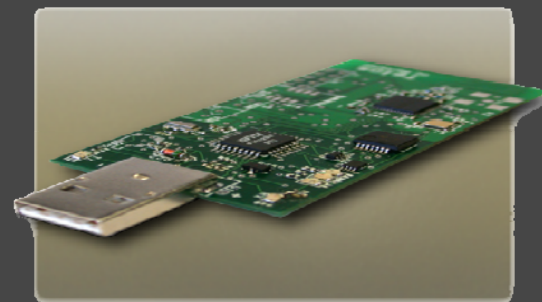

Telos Mote

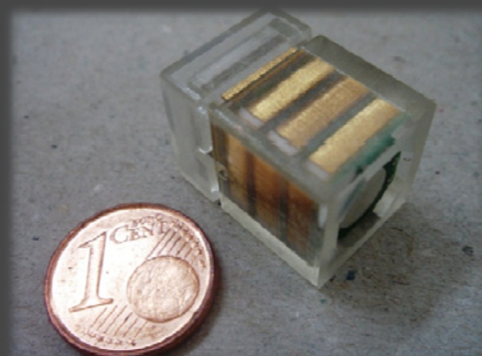

UCB PicoCube

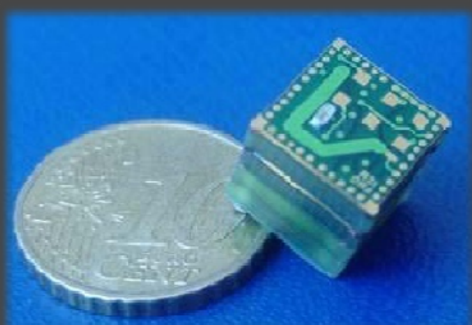

IIMEC e-Cube

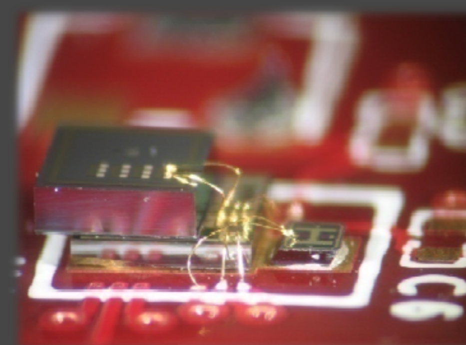

UCB mm^3 radio

[Ref: Ambient Intelligence, W. Weber Ed., 2005]

Yet … True Immersion Still Out of Reach

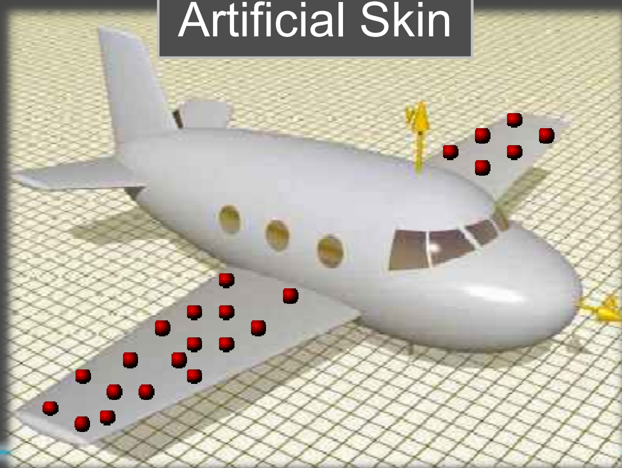
Artificial Skin


Interactive Surfaces


Smart Objects

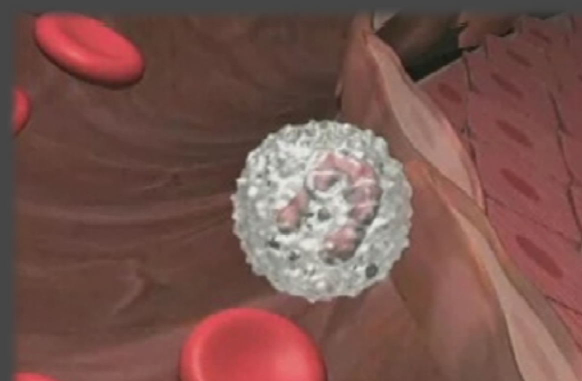
“Microscopic” Health Monitoring

Another leap in size, cost and energy reduction

Rethinking the Meaning of Scaling

- Traditional scaling rules to have minor impact
- Scaling is in the number of components, not in the transistor sizes
- A path to "More Than Moore" or "Beyond Moore"

"More Than Moore"

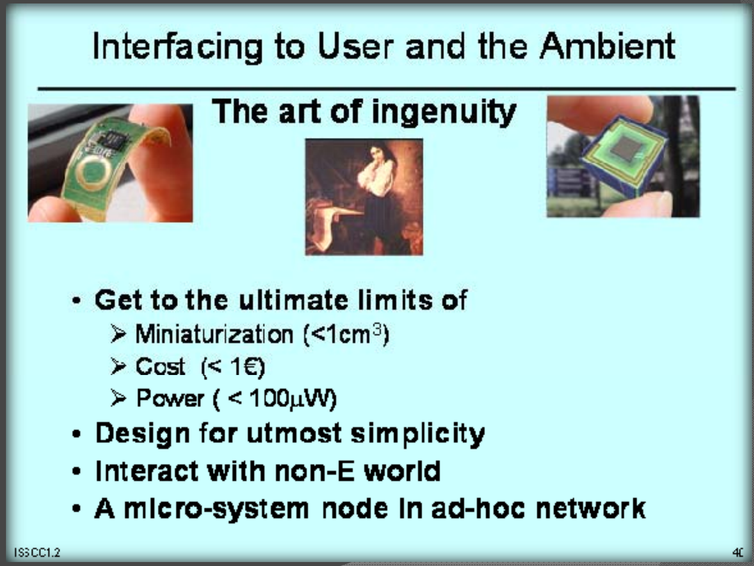

[H. De Man, Keynote Address, ISSCC 2005]

More Than Moore — Driven by Technology Innovation

Passive MEMS Components Provide Selectivity at ULP [Courtesy: N. Pletcher, UCB]

Nanowire-based AM Radio [Courtesy: Jensen, UCB]

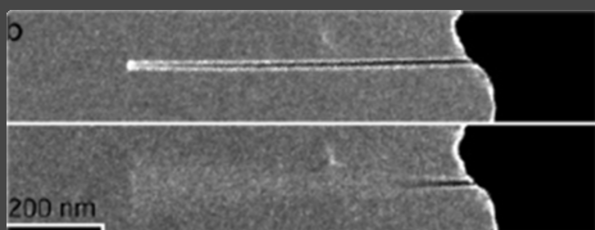

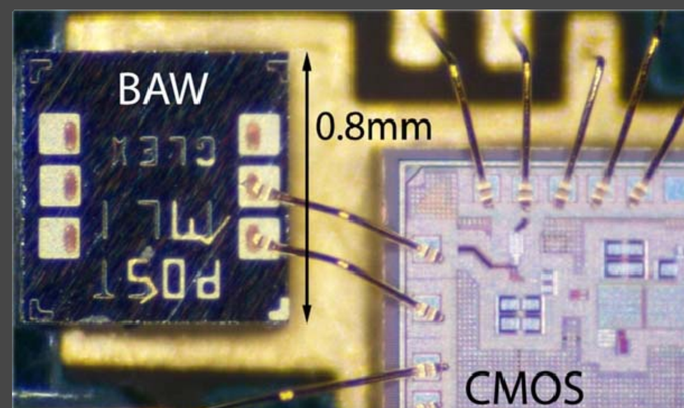

Mechanical Computing [Courtesy: C. Nguyen, UCB]

Beyond Moore …

- True immersion means broadening of the senses as well as "perceptional processing"
- "Bio-inspired" and "Bio-based" computing to lead to improved "user experience"

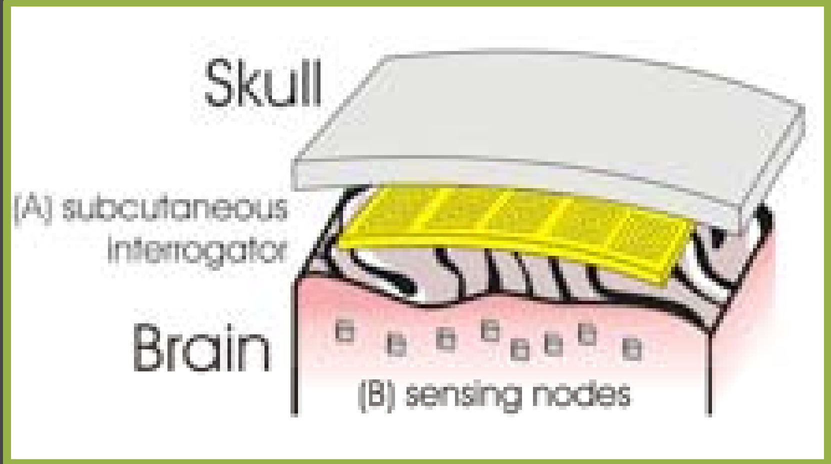

Immersed neural activity sensor: Potassium-modulated resonator [Courtesy: M. Maharbiz, UCB]

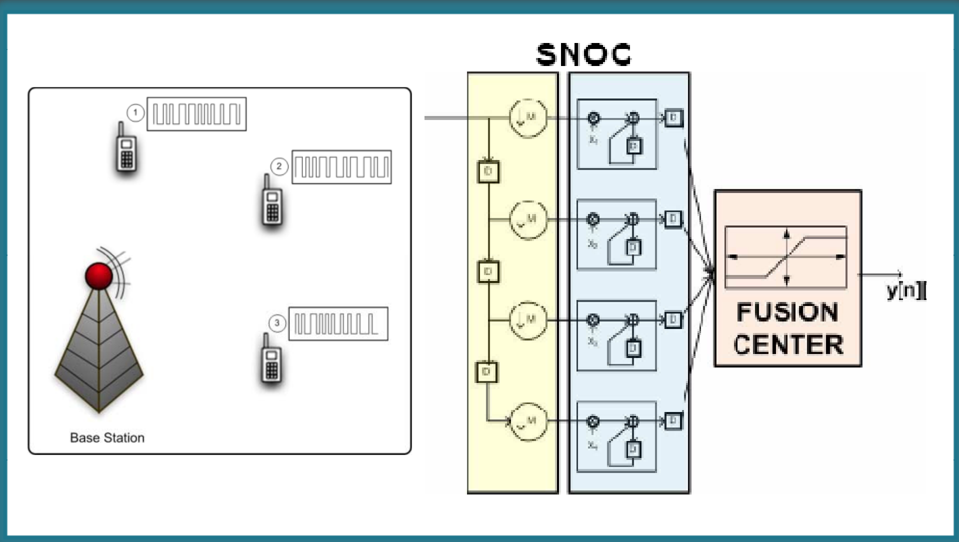

Sensor-network on a chip [Courtesy: N. Shanbhag, UIUC]

Yet …The TRUE Opportunity is in the TOTAL System

What Does This Mean for the Design Community?

The moment for true System-Level Design (SLD) is finally here

Quo Vadis, SLD? Reasoning About the Trends and Challenges of System Level Design

Recognizing common requirements for co-design of hardware and software in diverse systems may lead to productivity gains, lower costs and first-pass design success.

By Alberto Sangiovanni-Vincentelli, *Fellow IEEE*

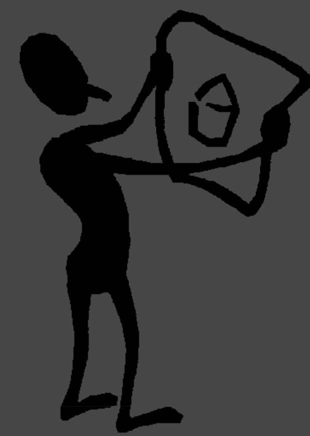

Proceedings of the IEEE, March 2007.

"… there is a common underlying basis that can be explored. This basis may yield a novel EDA industry and even a novel engineering field that could bring substantial productivity gains not only to the semiconductor industry but to all system industries including industrial and automotive, communication and computing, avionics and building automation, space and agriculture, and health and security, in short, a real technical renaissance." - ASV

A New Meaning to "System Design"

- Semiconductor and design automation industries focused on "component design"
- Need to address the system in a holistic way

Complexity and emergent behavior of networked systems

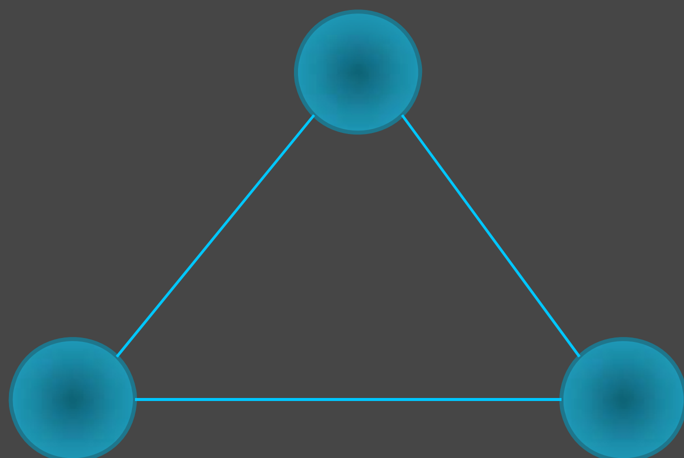

System-level reliability and liability

System-level metrics (energy, latency)

Addressing Complexity

- Directly caused by massive concurrency and heterogeneity in SiS systems

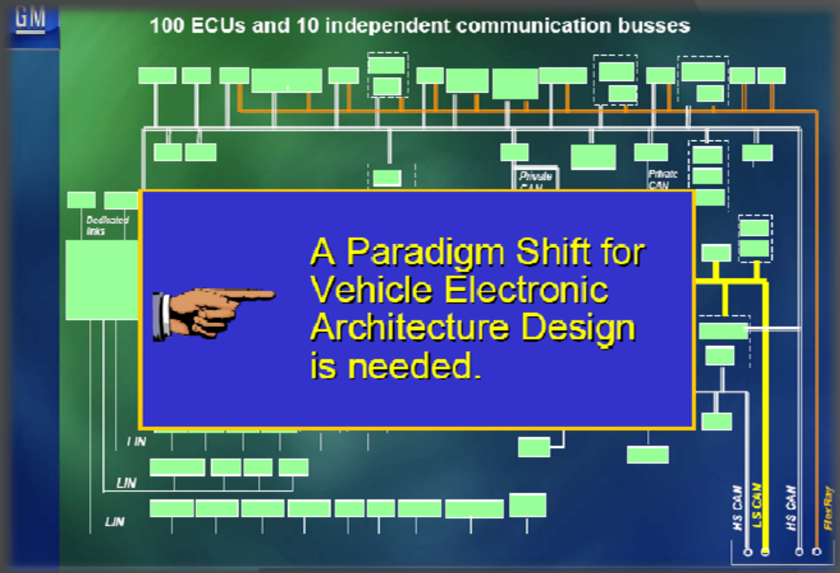

[Courtesy: M. Osella, GM]

- What is needed:
  - Raising the abstraction model
    - Fundamental change in existing “bottom-up” (“up-integration”) business model
  - Enabling a “virtual engineering” design methodology
  - A system-level design science

Addressing Reliability

- A system-level responsibility
  - Reliability can and should not be provided by components alone
  - Correct system behavior does not require determinism at all levels
- Reliability modeling requires statistics
  - Models and abstractions that express reliability requirements and capabilities at all hierarchy levels
  - Needs revisiting of verification strategies

Addressing Reliability

- Redundancy and resiliency the essential tenets
  - Exploit the “swarm” component of SiS

10-15% of terrestrial animal biomass

10^9 Neurons/”node”

Since 10^5 years ago

10-15% of terrestrial animal biomass

10^5 Neurons/”node”

Since 10^8 years ago

Easier to make ants than humans

“Small, simple, swarm”

Revisiting the Metrics

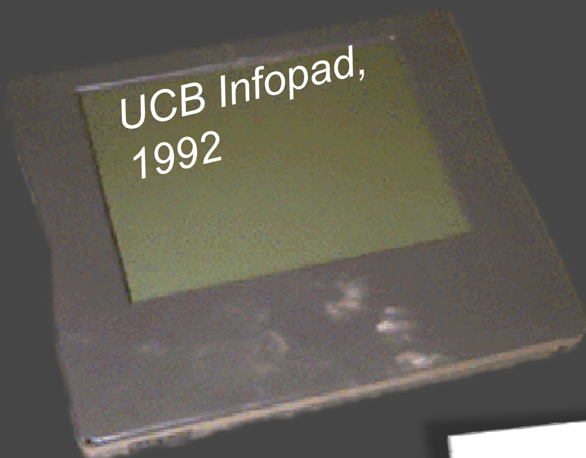

- In “always-connected” world, energy-intensive tasks can be performed in “power-rich” backbone
  - Use energy when and where available

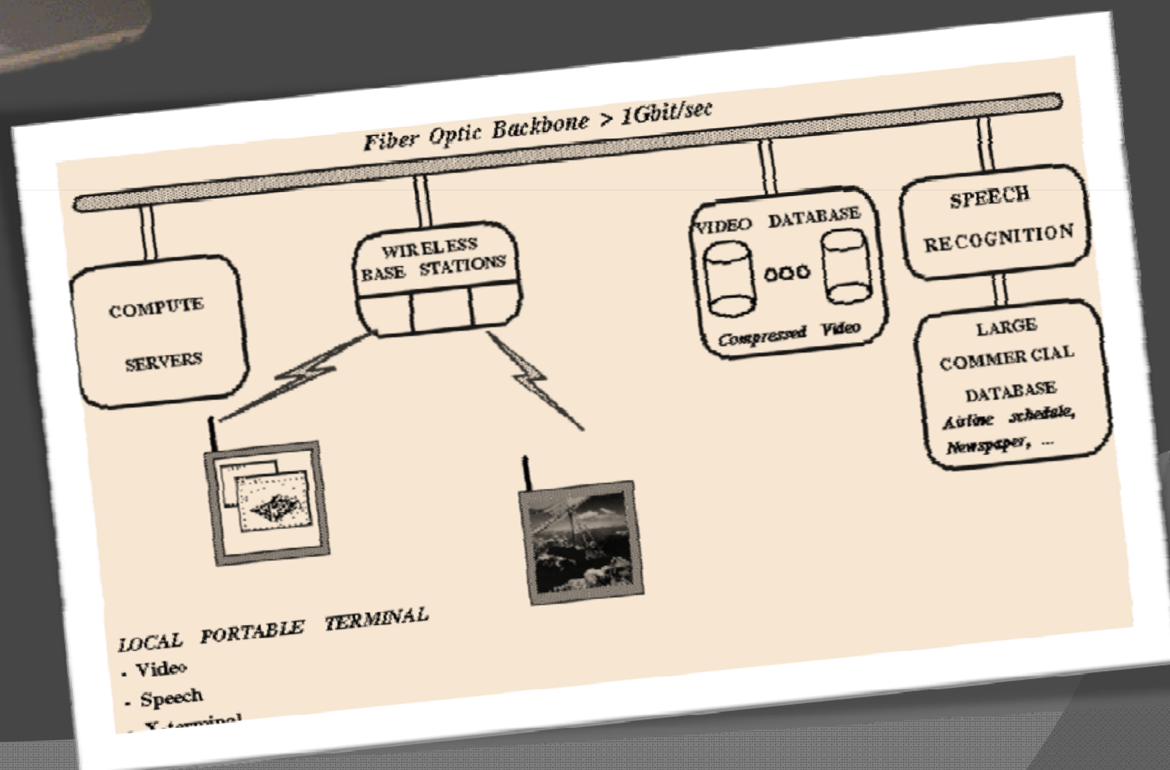

Energy and Latency as Dominant Metrics

Energy efficiency and performance as a network problem – not so much a device issue!

- What matters is "perceived user experience/unit energy consumed" (*)
  - Other interesting metrics: lifetime of network given energy available at nodes
- Requires trade-off between cost of computation and communication, as well as overcoming latency constraints
  - Latency matters more than performance

(*) Term first coined by John Shen (Nokia NRC)

Fractal Nature of Design:

Platform-Based Design Covers all Scales

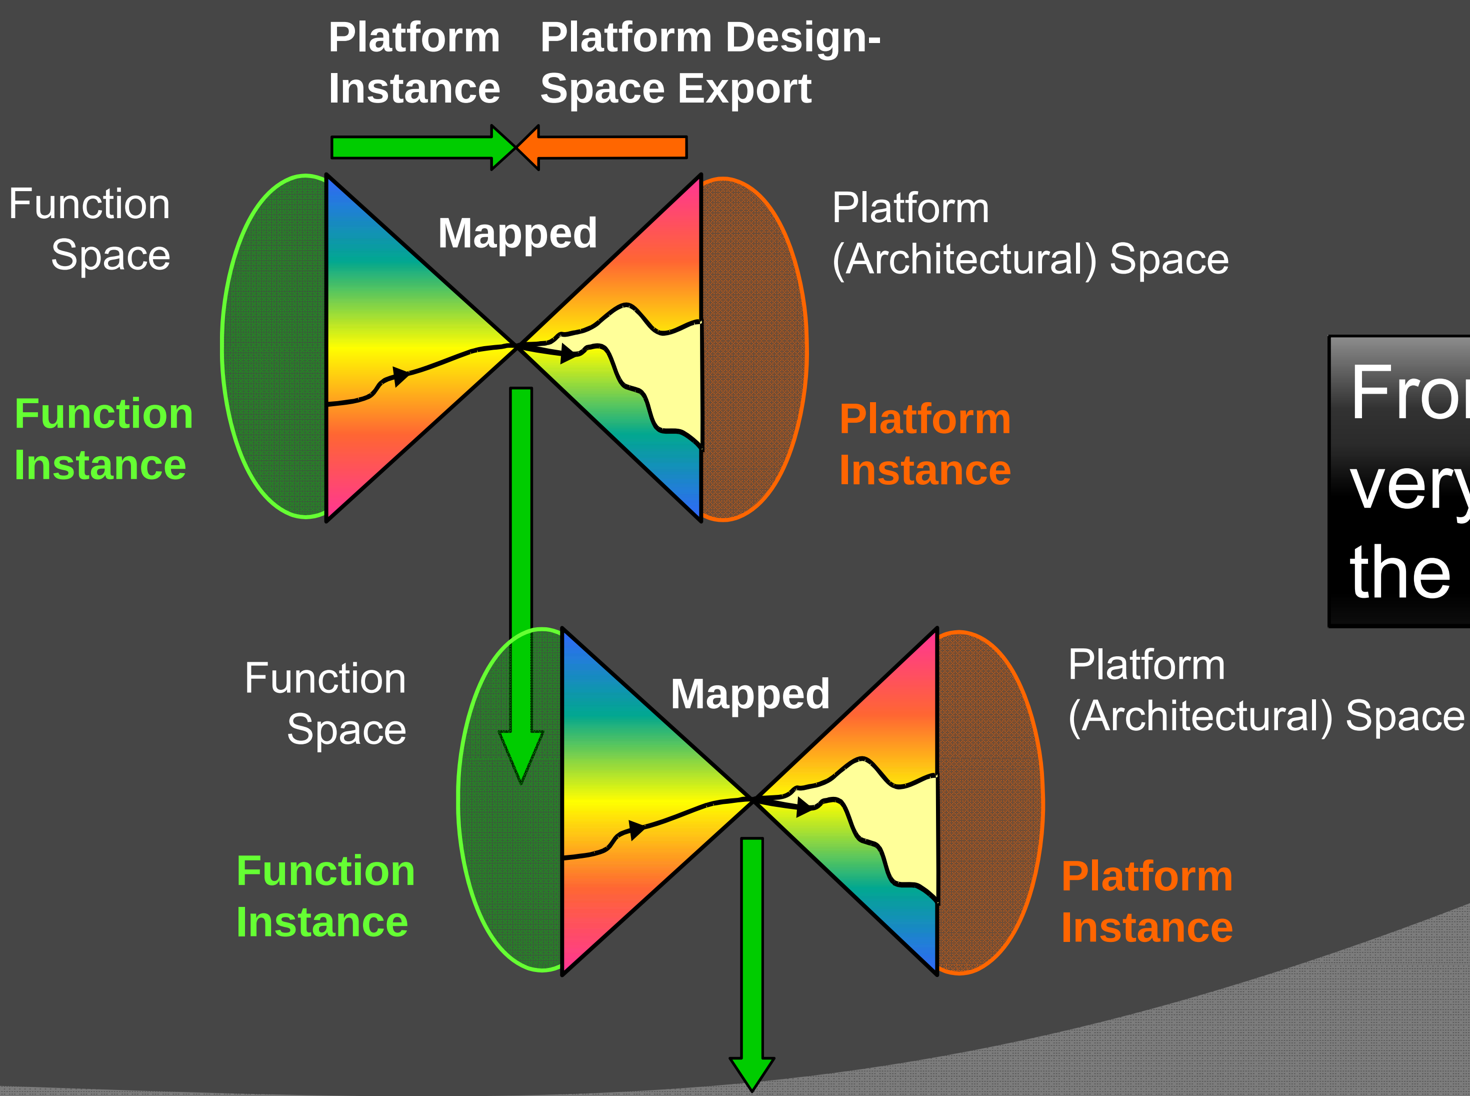

From the very large to the very small

[Courtesy: A. Sangiovanni-Vincentelli]

Major Take-Away's

- Ubiquitous always-connected wireless radically transforming the Information Technology Arena
  - Towards truly Immersive Systems
- Complexity, reliability and energy present formidable challenges
- "Design Technology" has to extend itself from "component" to system oriented
  - Must subsume traditional design flows rather than replace them
- Broad collaboration between systems and semiconductor industries, as well as industry and academia needed
  - Need for new benchmark libraries
  - Need theory of system design

[J. Rabaey et al., "Workloads of the Future", IEEE D&T Magazine, July/August 2008]