Reliability in the Dark Silicon Era

Babak Falsafi
Director, EcoCloud
ecocloud.ch

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

IT is ever more **indispensable**

Our life w/o digital data is unimaginable as

- Enterprises
- Scientists
- Governments
- Societies
- Individuals

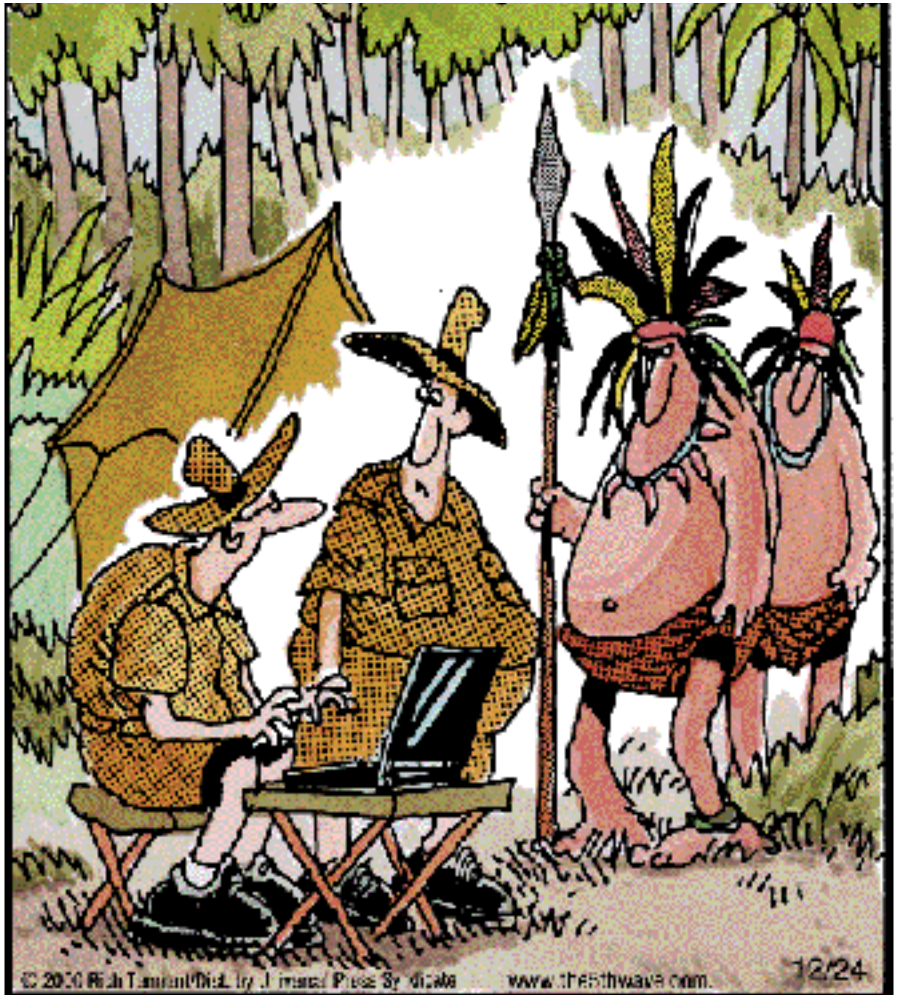

"He saw your laptop and wants to know if he can check his Hotmail."

IT: An Exponential Growth

Intel 4004, 1971

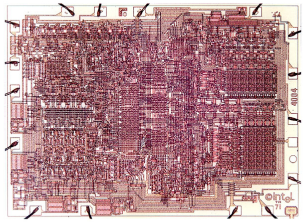

92,000 ops/second

Intel Nehalem, 2009

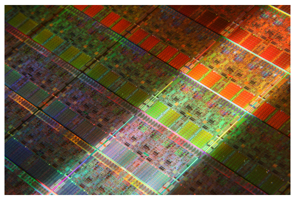

12,000,000,000 ops/second

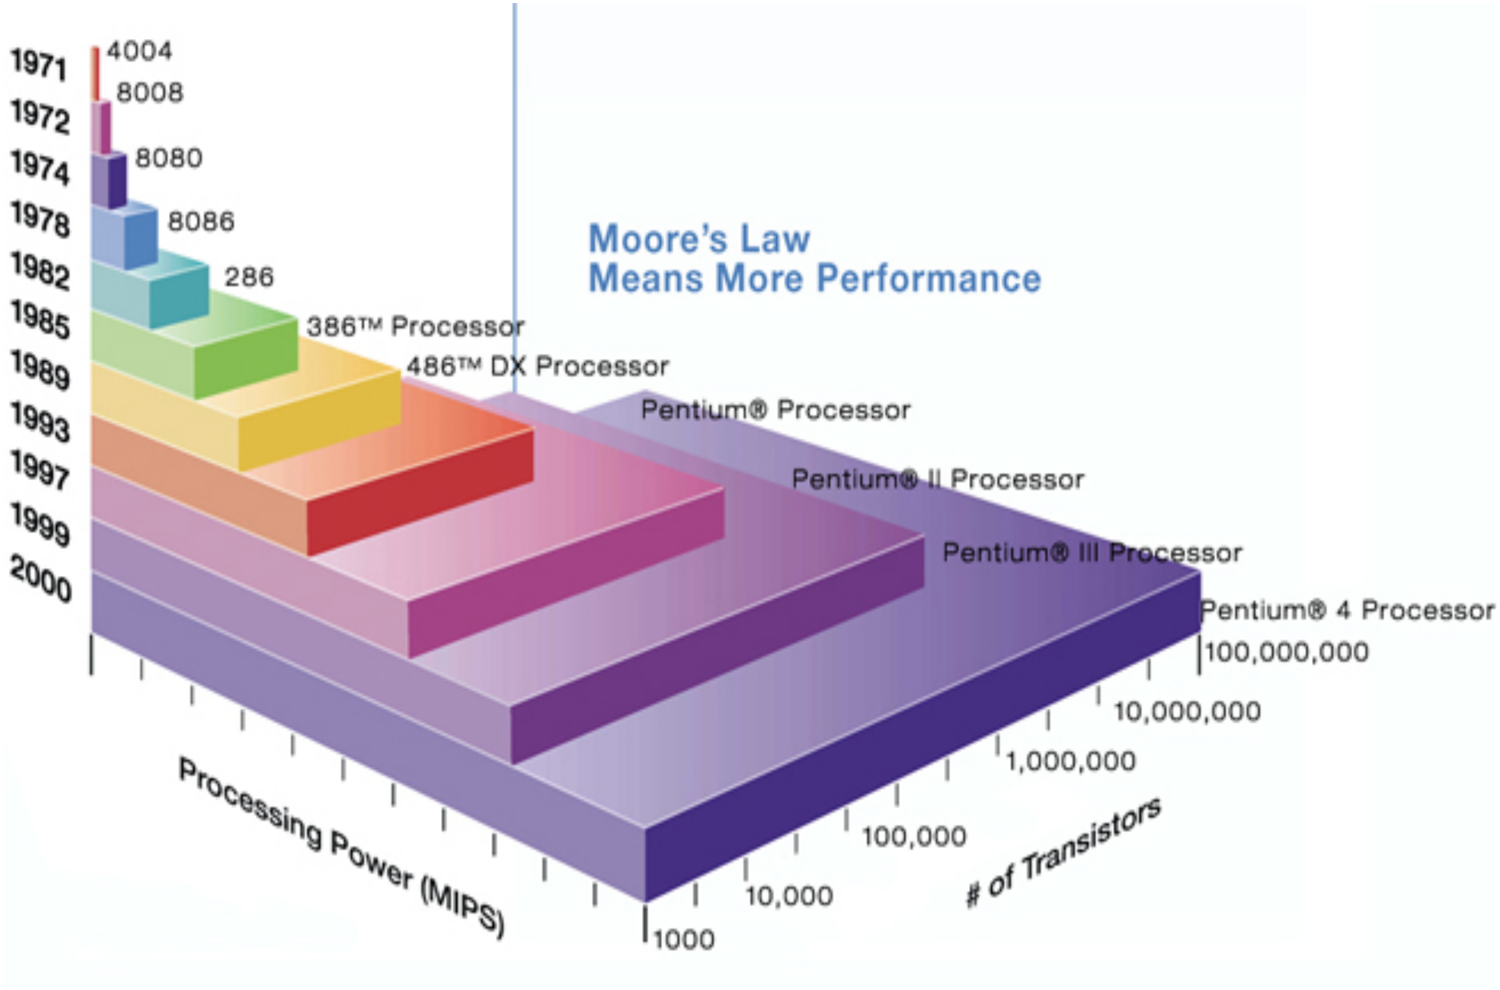

Four decades of digital platform proliferation

Exponential increase in density & decrease in cost

A Brief History of IT

Communication Era

Consumer Era

1970s- 1980s 1990s Today+

Mainframes

PC Era

- From scientific instrument to commodity
- From product to service

IT: The Consumer Era

Phenomenal change from decades ago:

- Instant connectivity
- Shopping now online
- Daily interaction > 300 people
- Augmented reality
- Streaming movies
- …..

IT is at core of everyone's life!

Change in IT's Landscape

➔ Emergence of Data-Centric Universe
  – IT focus on massive data

- End of Dennard Scaling
  – Higher density ➔ higher energy

- Data-Centric Universe meets Energy Wall

What are design implications?

Our Data-Centric Universe: Data Growing faster than Technology

- Commerce entirely data-driven
- Science handling massive data
- Companies spending $$$ to collect/analyze data
- Personalized computing

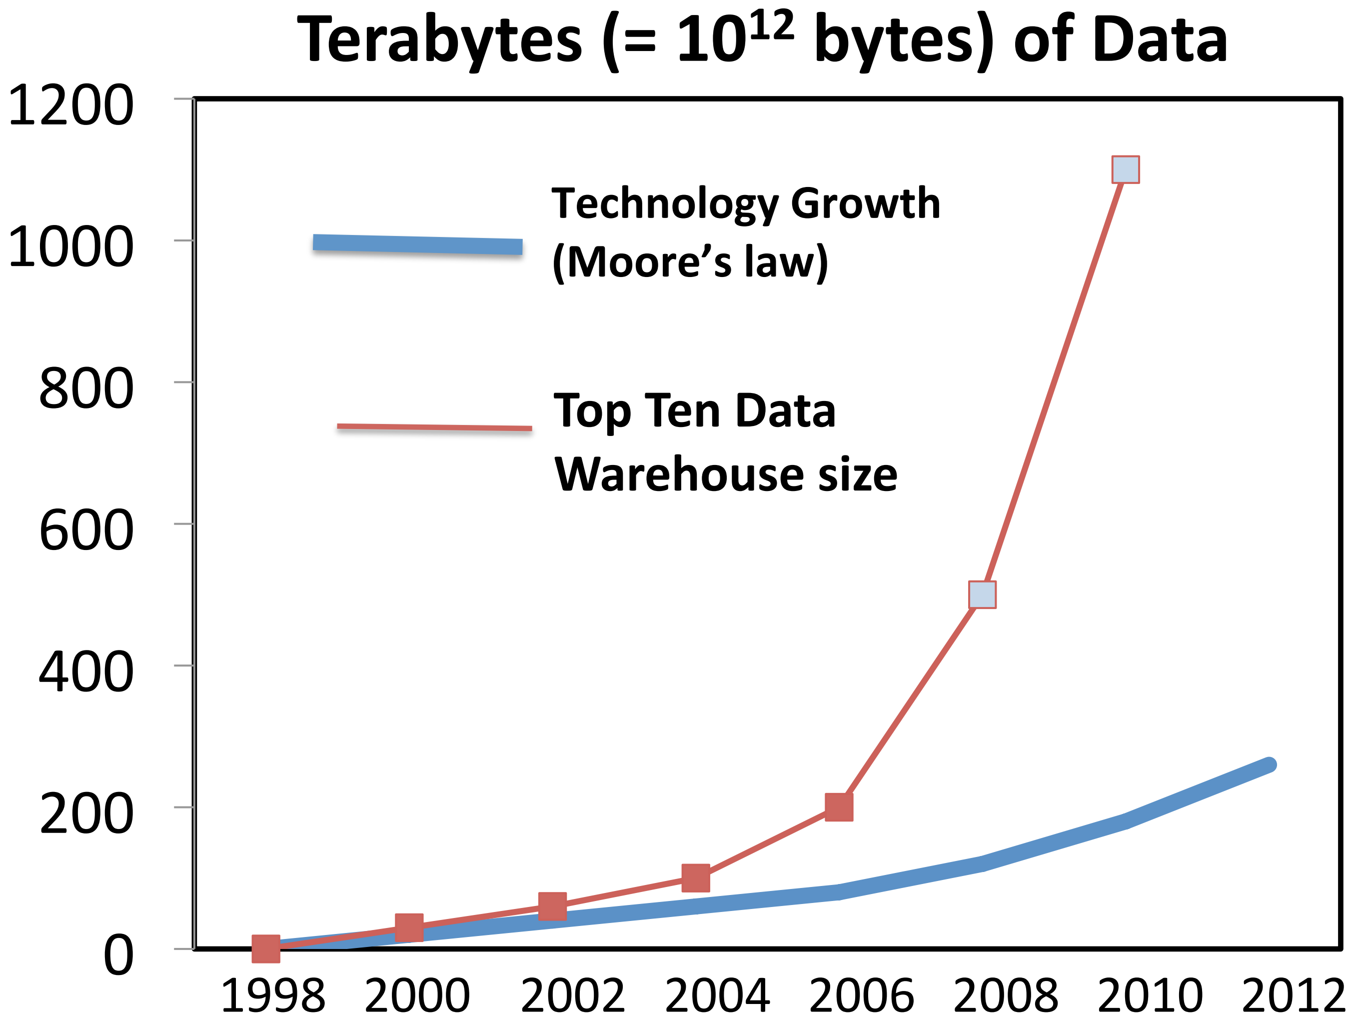

WinterCorp Survey, www.wintercorp.com

Data Deluge: 1200 Exabytes in 2010

(Economist, Feb. 25th 2010)

- Only 150 Exabytes in 2005
- Supply-chain management, 10x increase in data in a year
- US aerial surveillance models 30x more data in 2011

Anatomy of a Data-Centric Business

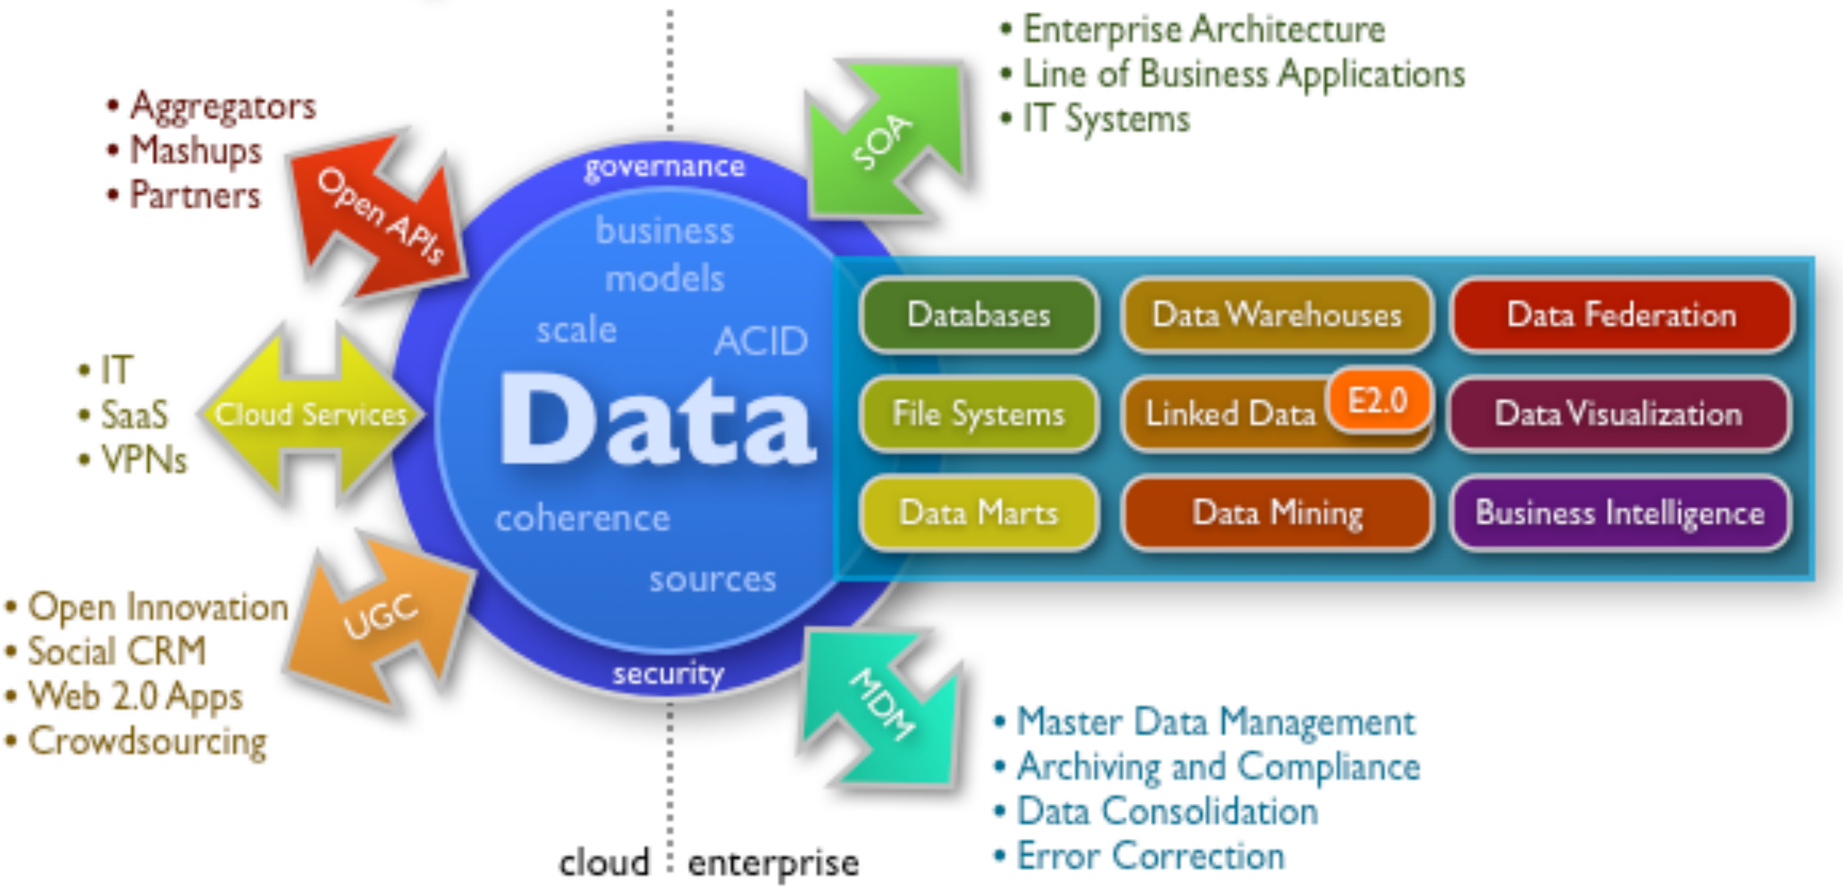

- Era of “knowledge economy”
- 50% of economic value in developed countries
- Dominant supply-chain component of products/services

Data-Centric Science: “The Fourth Paradigm”

Mining data from:

- Archives
- Humans
- Sensors/instruments
- Simulations

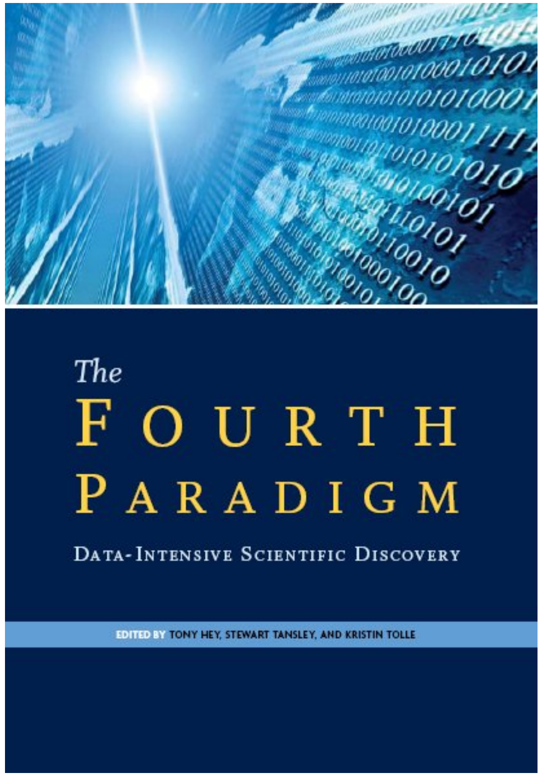

Unifying theory, experimentation, simulation, analytics on massive data

Data Comes in Various Flavors

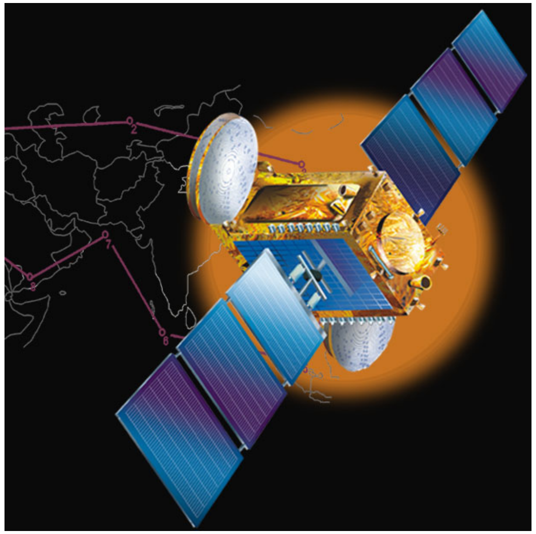

Satellite

Health

Entertainment

Life

Commerce

Search

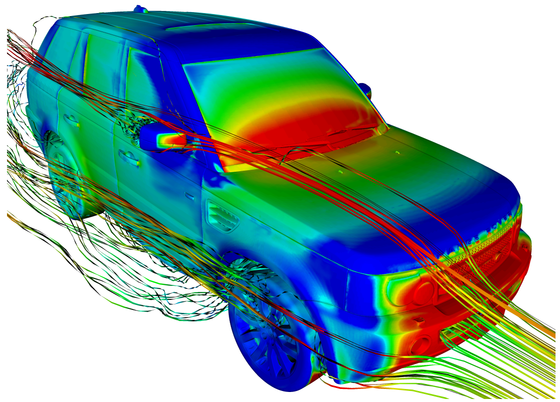

Simulation

It’s all about Accessing Data!

Cloud Computing

A computing paradigm shift to enable ubiquitous connectivity

How to design for massive data

Change in IT's Landscape

- Emergence of Digital Universe
  – IT focus on massive data

➔End of "Free Energy"
  – Higher density ➔ higher energy

- Data-centric Universe meets Energy Wall

What are design implications?

IT Energy is Shooting Up!

IT riding on technology that was energy-friendly

- Exponentially better performance, density
- Constant power envelope

But, energy is shooting up!

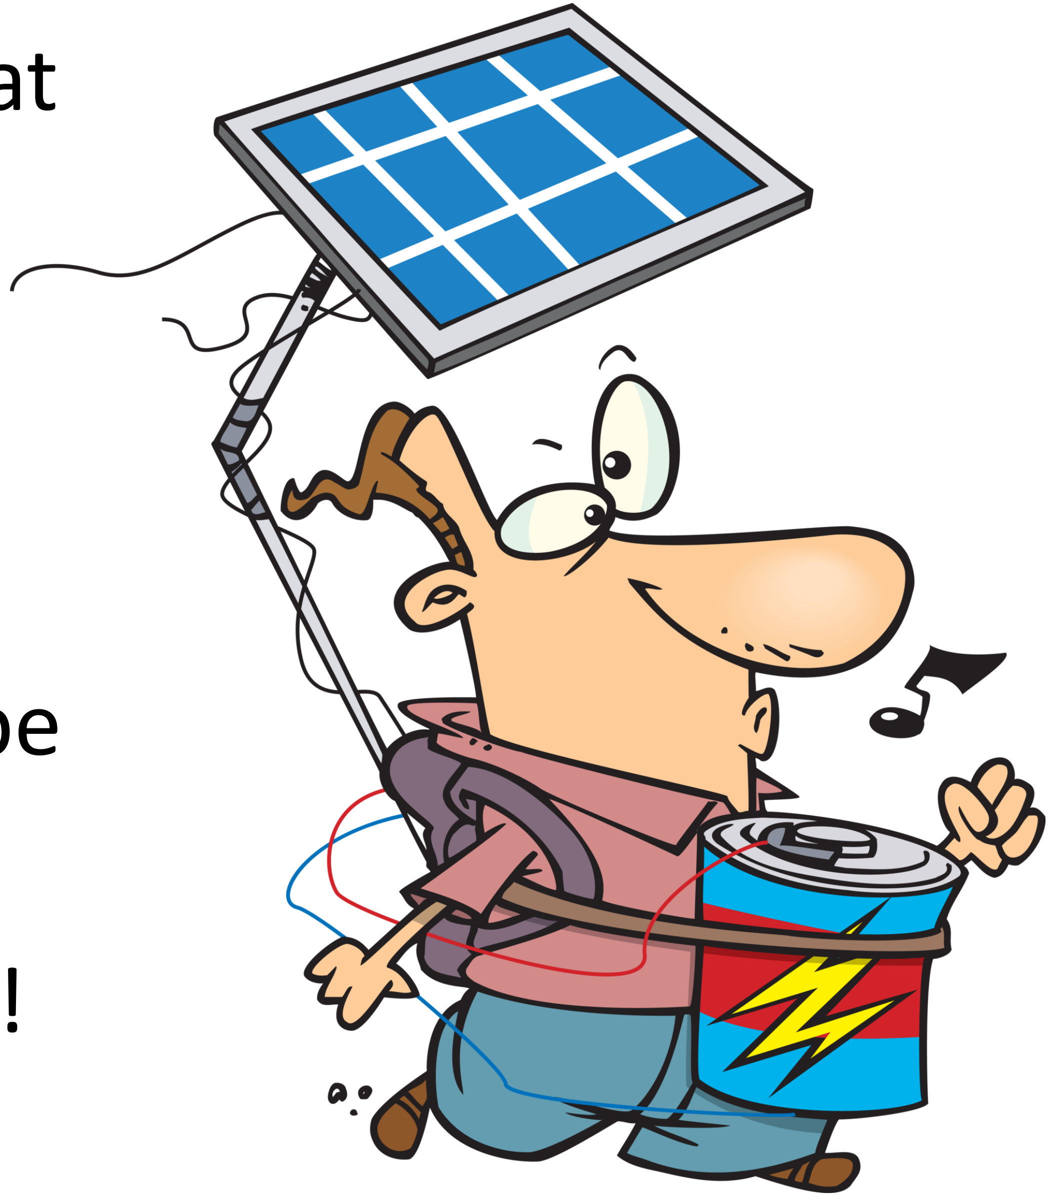

Household Energy in the UK
(UK BERR, 2008)

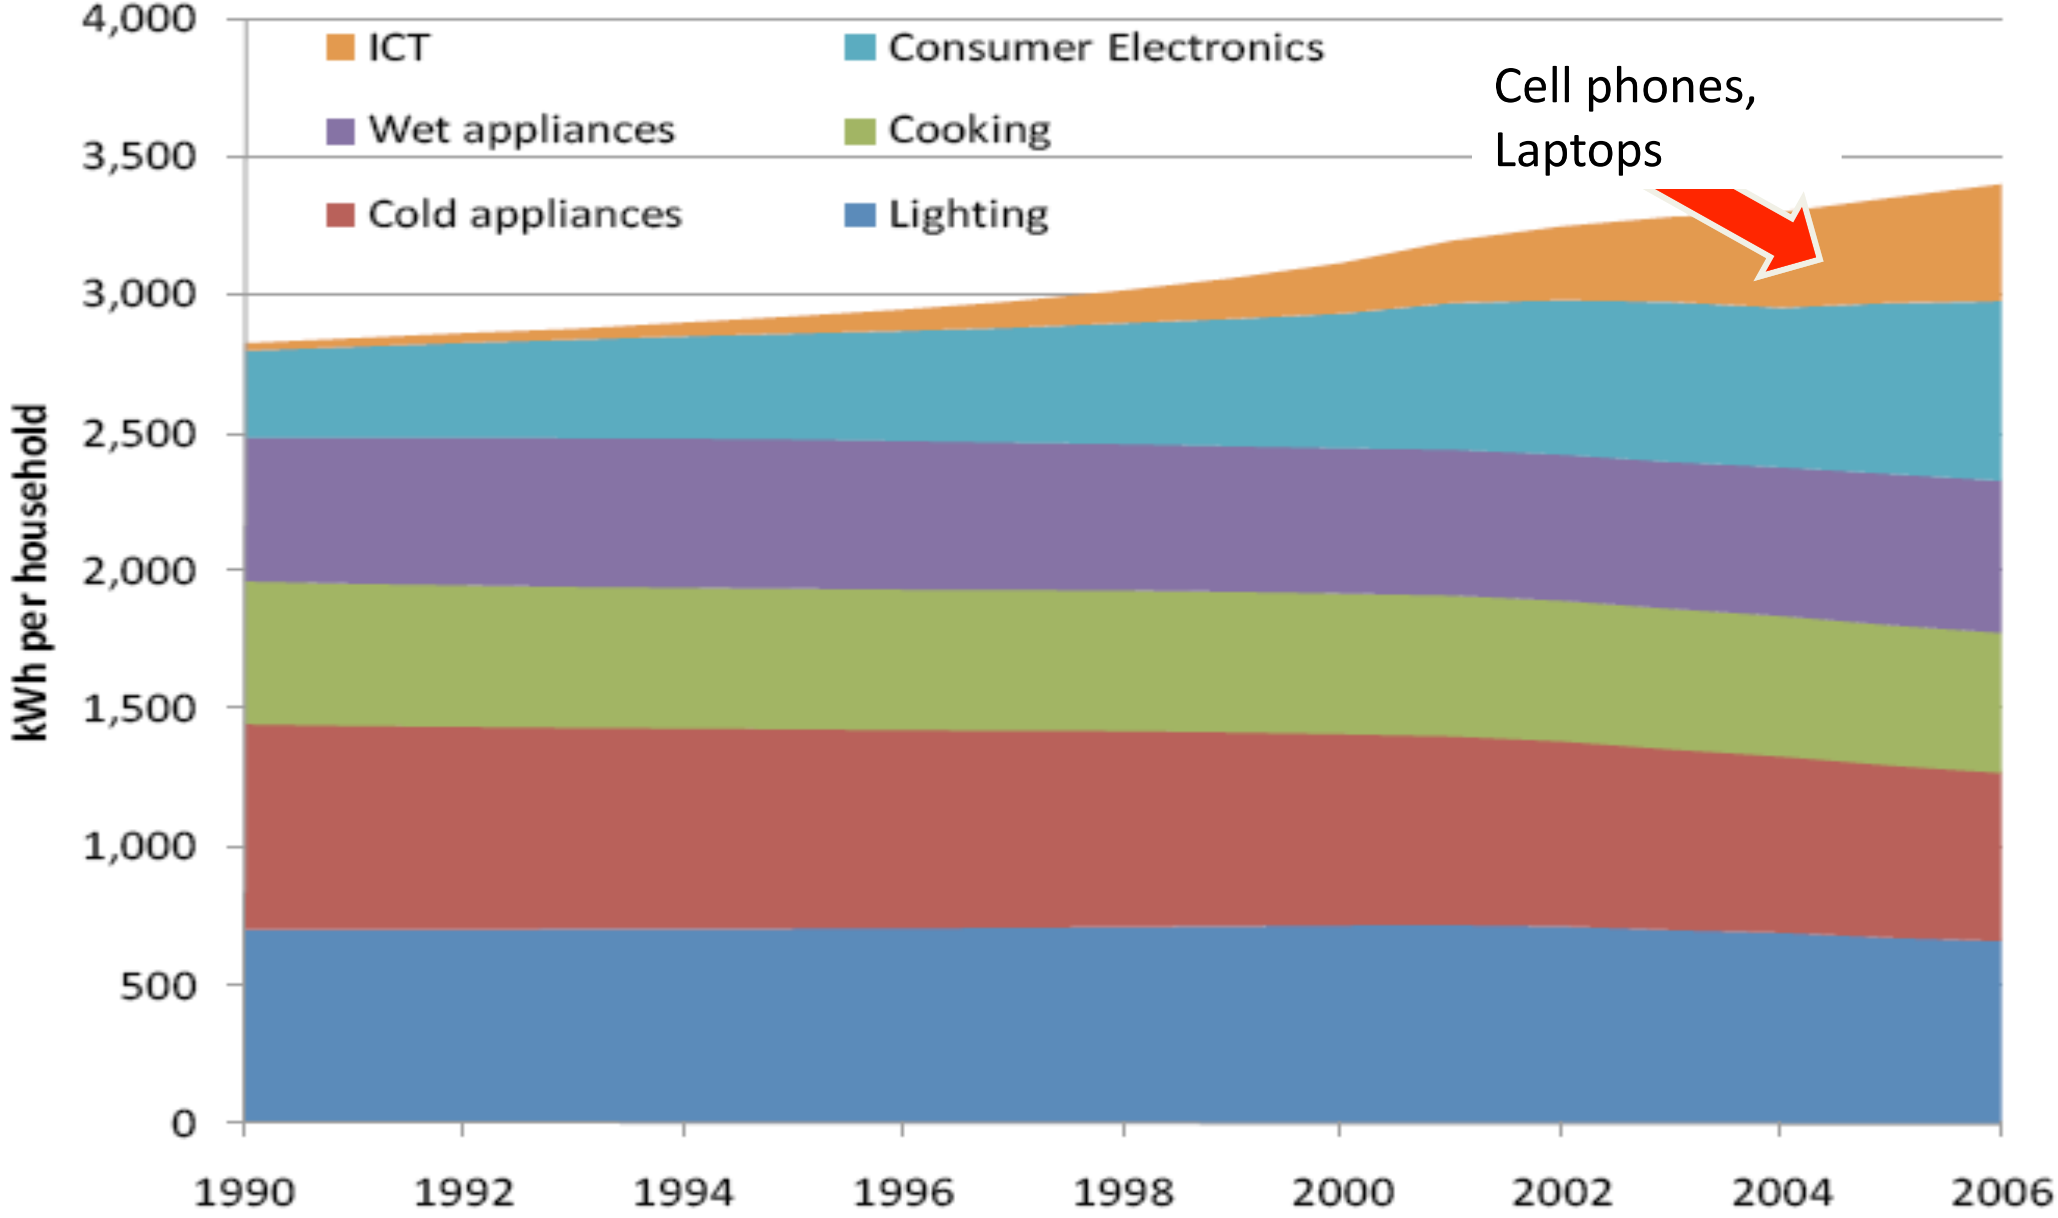

Household Energy in the US

(NY Times, 2011)

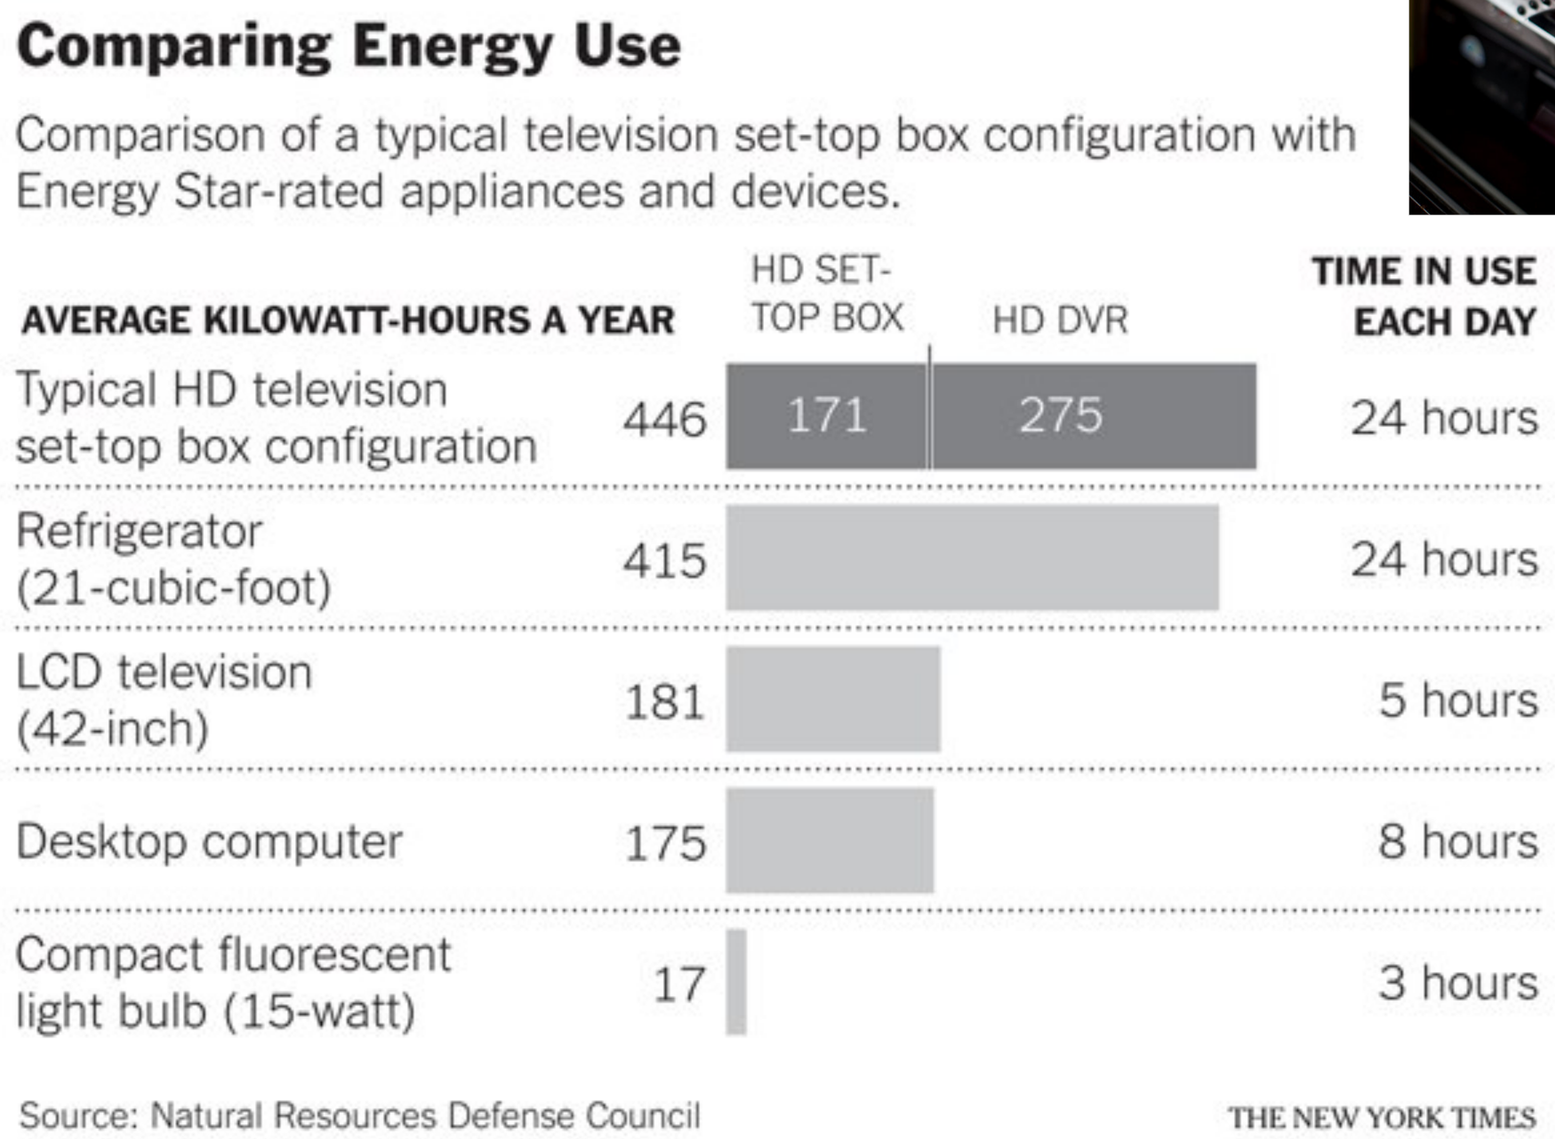

Data center Energy in the US

(extrapolated from Energy Star, 2007)

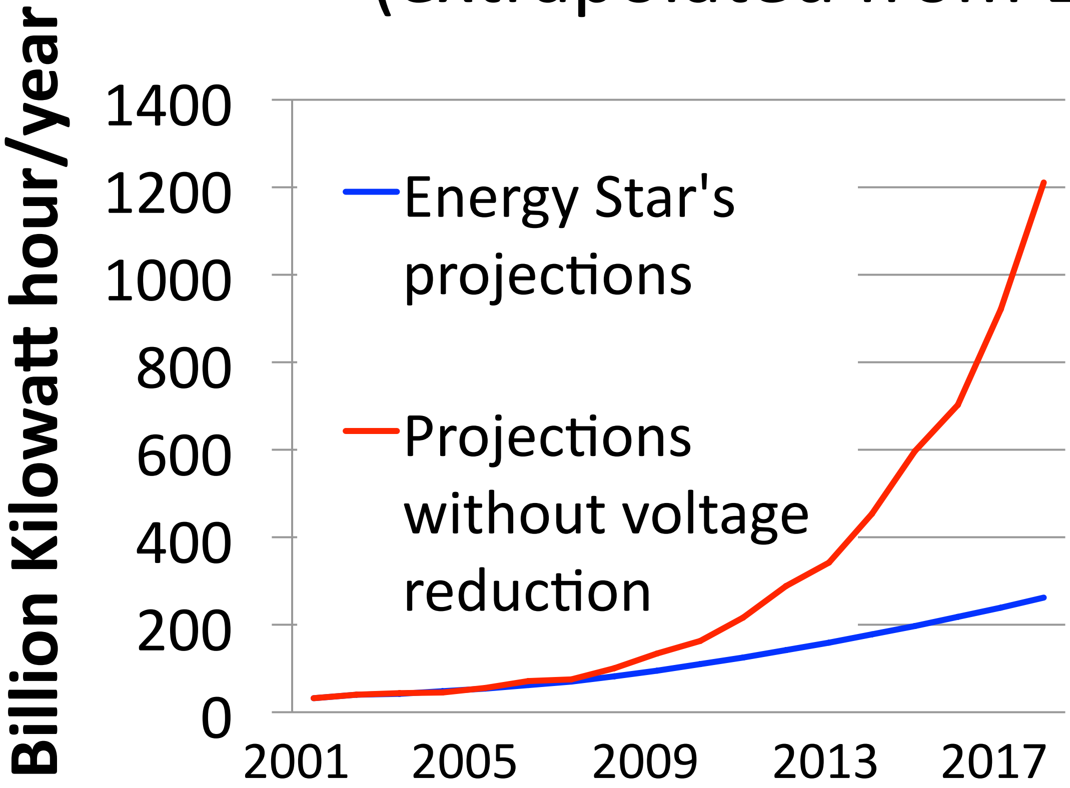

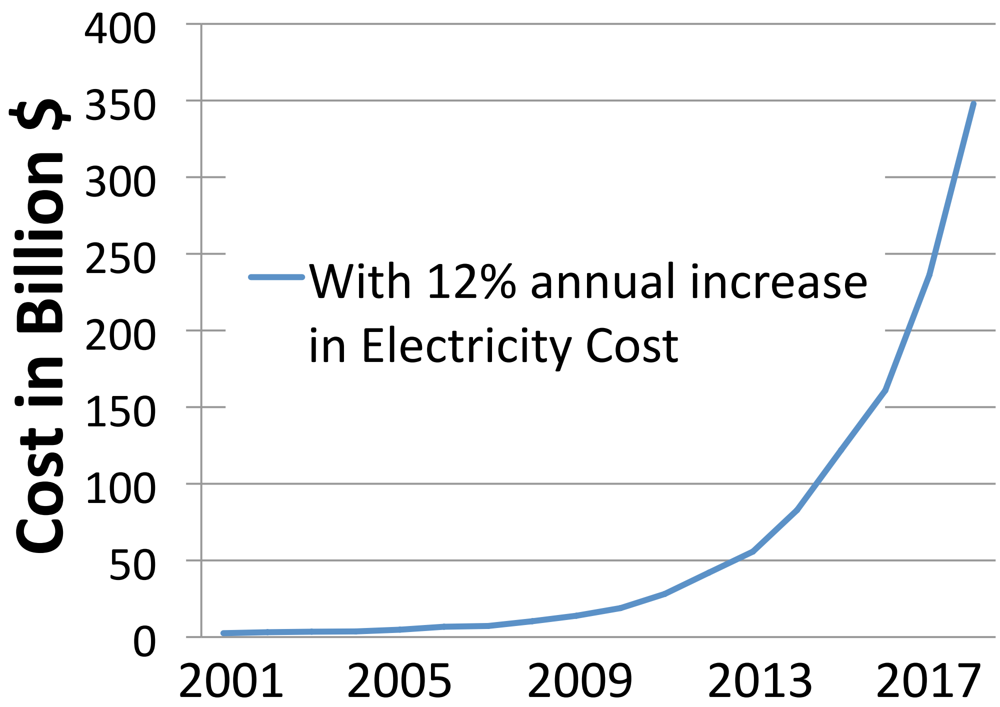

- Exponential costs if not mitigated
- Today, carbon footprint of airline industry

Energy > Capital Cost

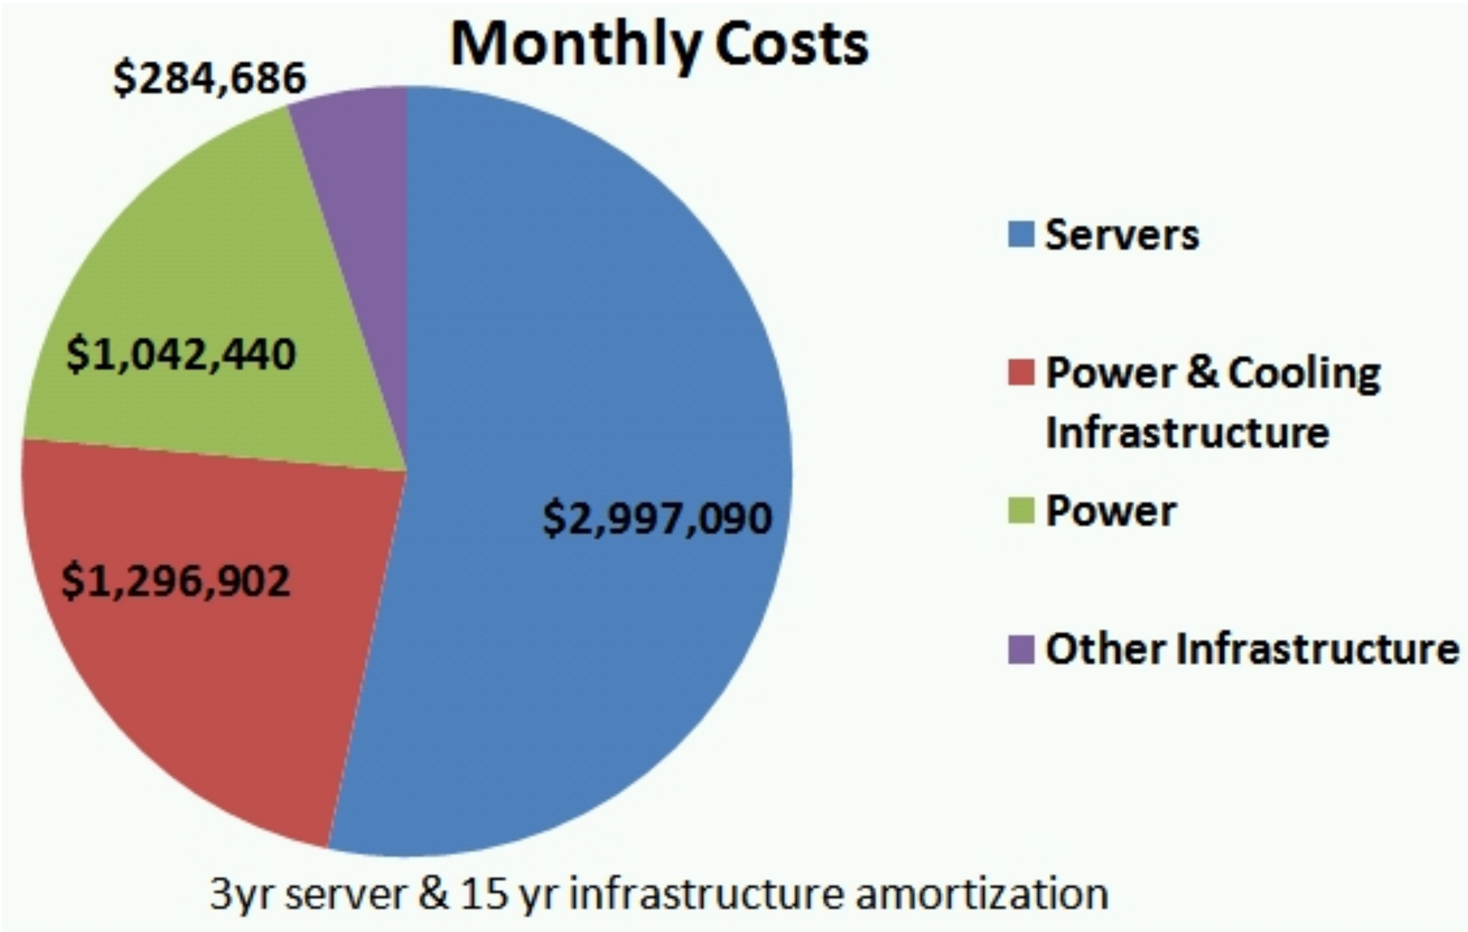

James Hamilton's Blog, mvdirona.com, 2008

- Servers are getting relatively cheaper
- Power is beginning to dominate cost

Technology Inflection Point in 2004!

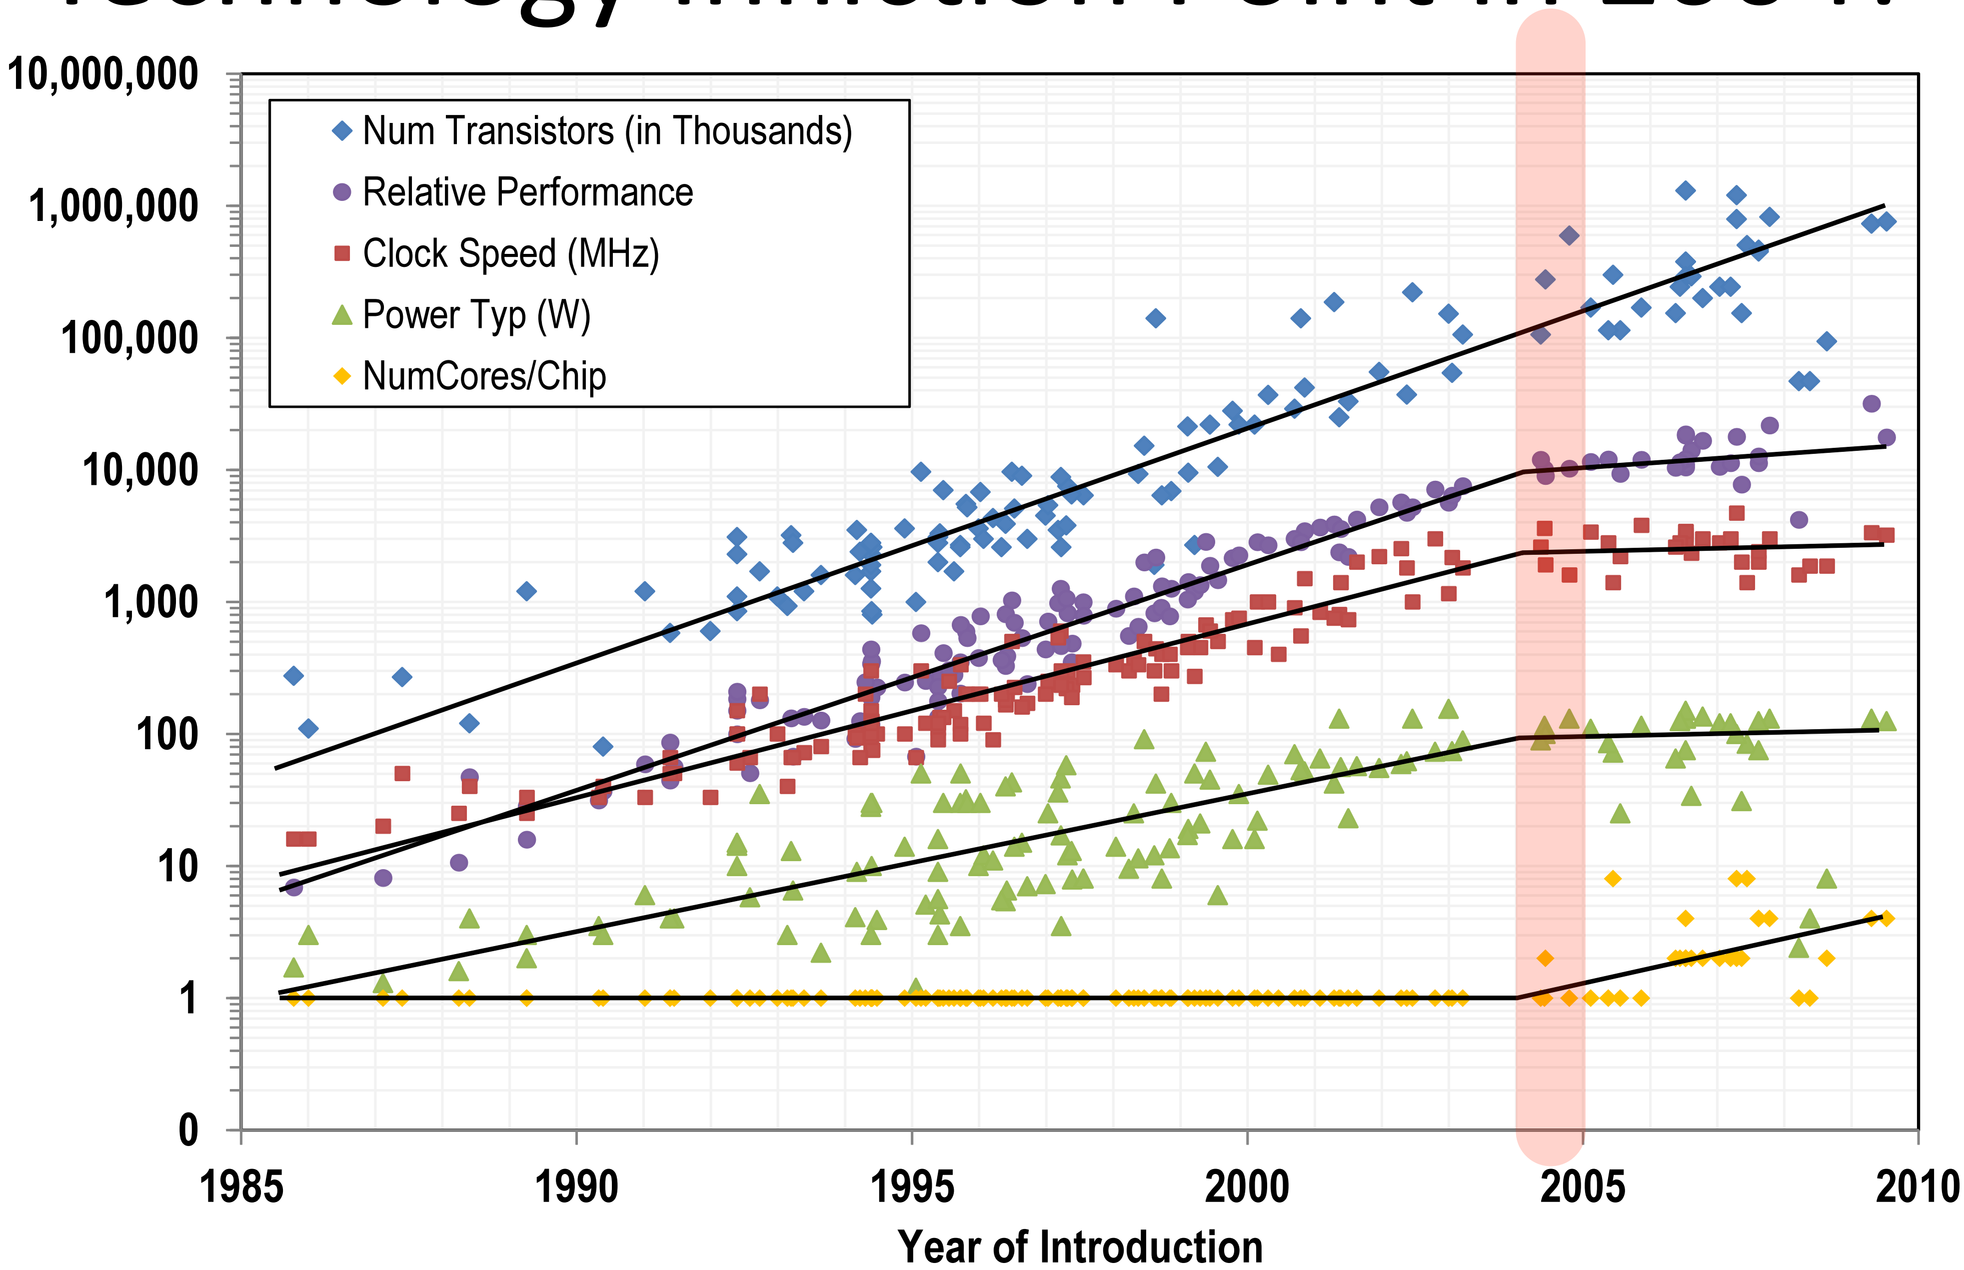

National Research Council (NRC) – Computer Science and Telecommunications Board (CSTB.org)

Four decades of Dennard Scaling

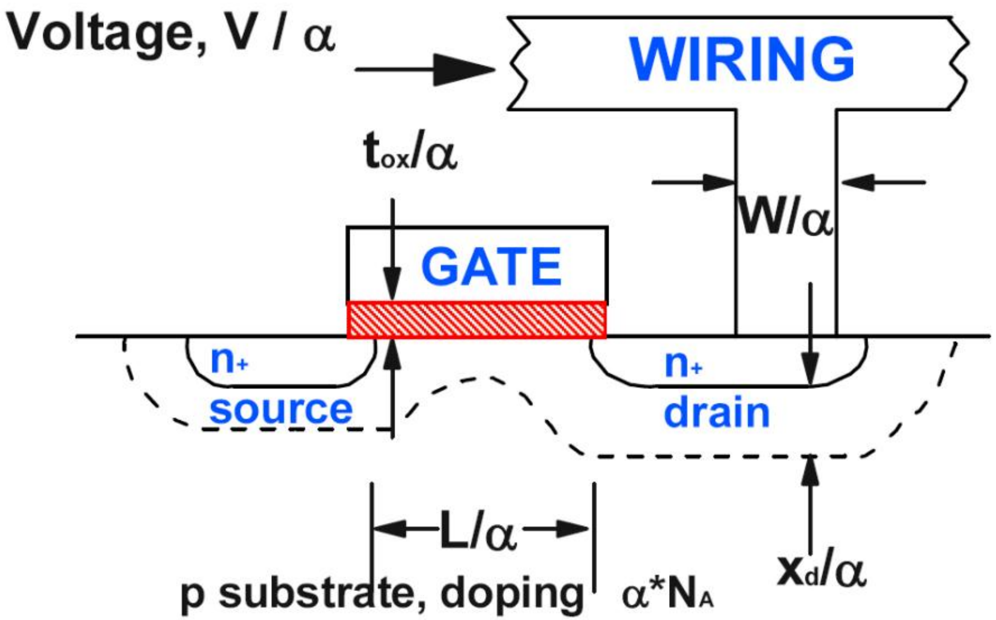

Dennard et. al., 1974

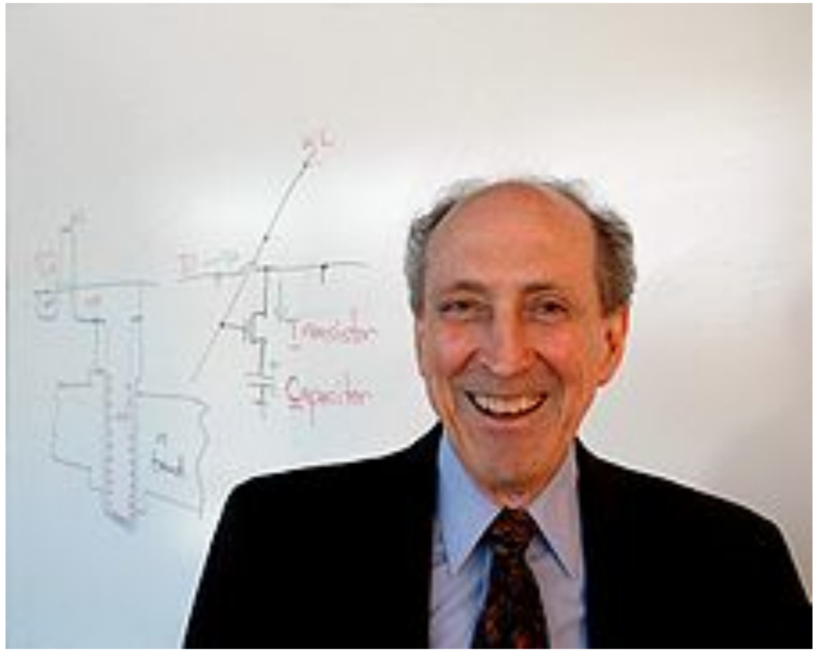

Robert H. Dennard, picture from Wikipedia

- **$P = C V^2 f$**
- Increase in device count
- Lower supply voltages

➜ Constant power/chip

Leakage Killed Dennard Scaling

Leakage:

- Exponential in inverse of V_{th}
- Exponential in temperature
- Linear in device count

To switch well

- must keep $V_{dd}/V_{th} > 3$

➔V_{dd} can't go down

End of Dennard Scaling (ITRS)

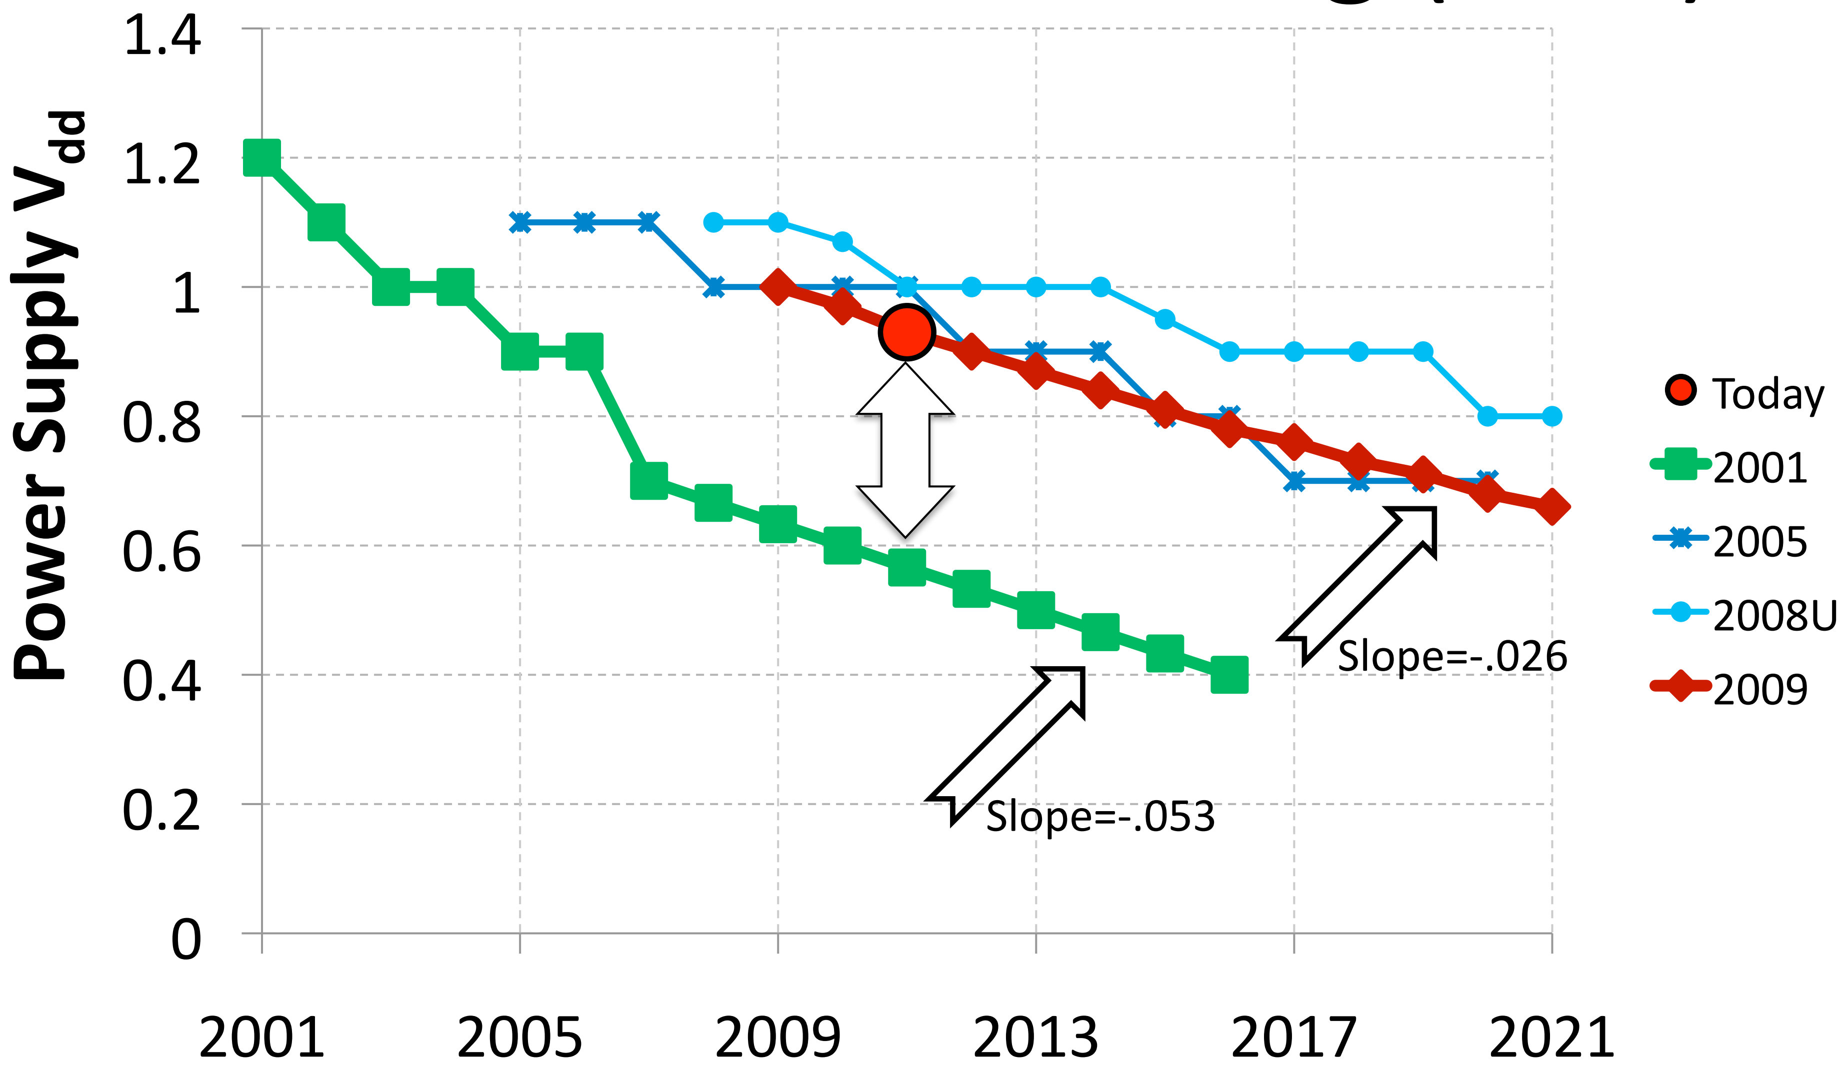

Mike Ferdman, from ITRS pages, July 2011

Supply voltages going down at much lower rate!

Dark Silicon: End of Manycore Scaling

Can not power up chip

Parallelism has limits
even in Servers!

Must:

- specialize
- selectively power up

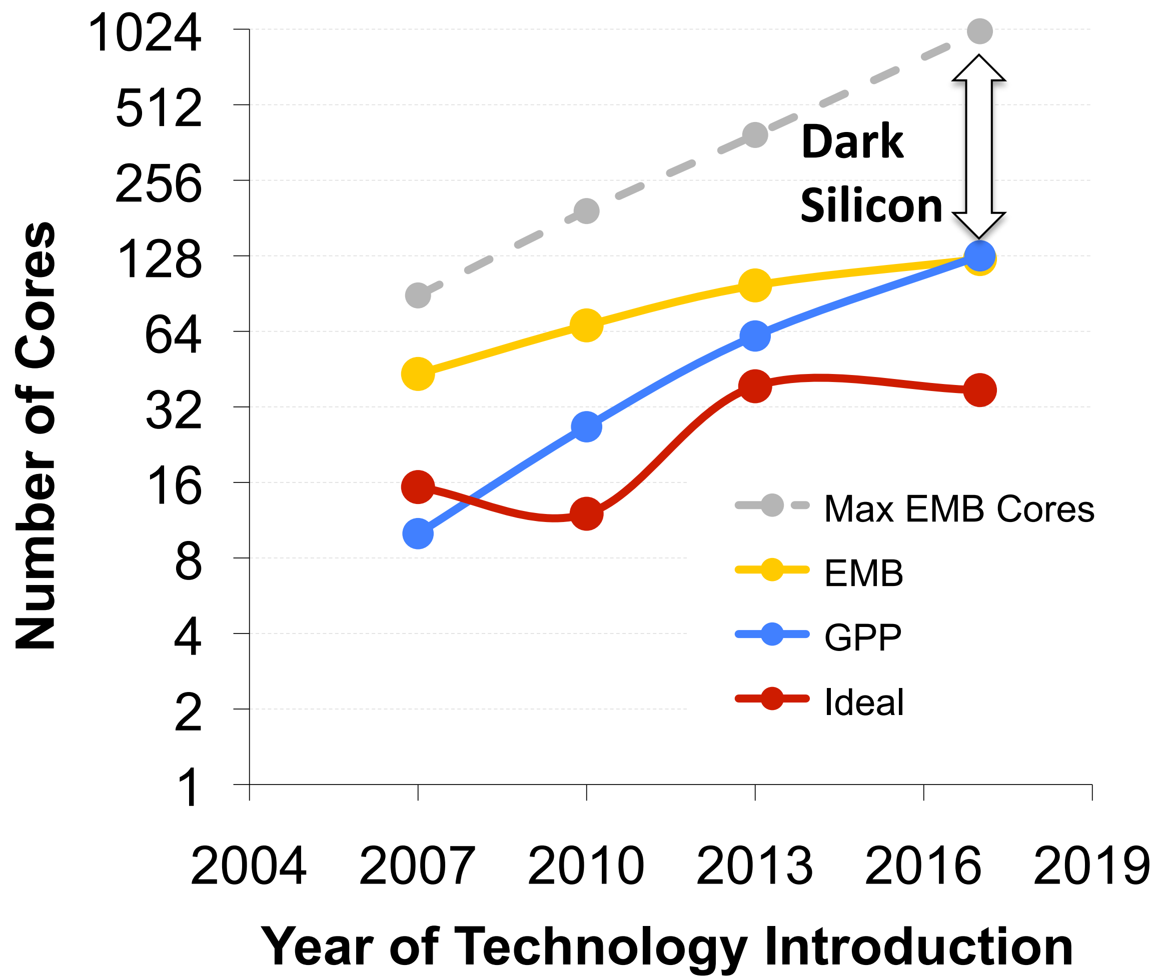

Hardavellas et. al., “Towards Dark Silicon in Servers”, 2011

Specialization can buy 1000x in Energy
(from a sample of 20 chips)

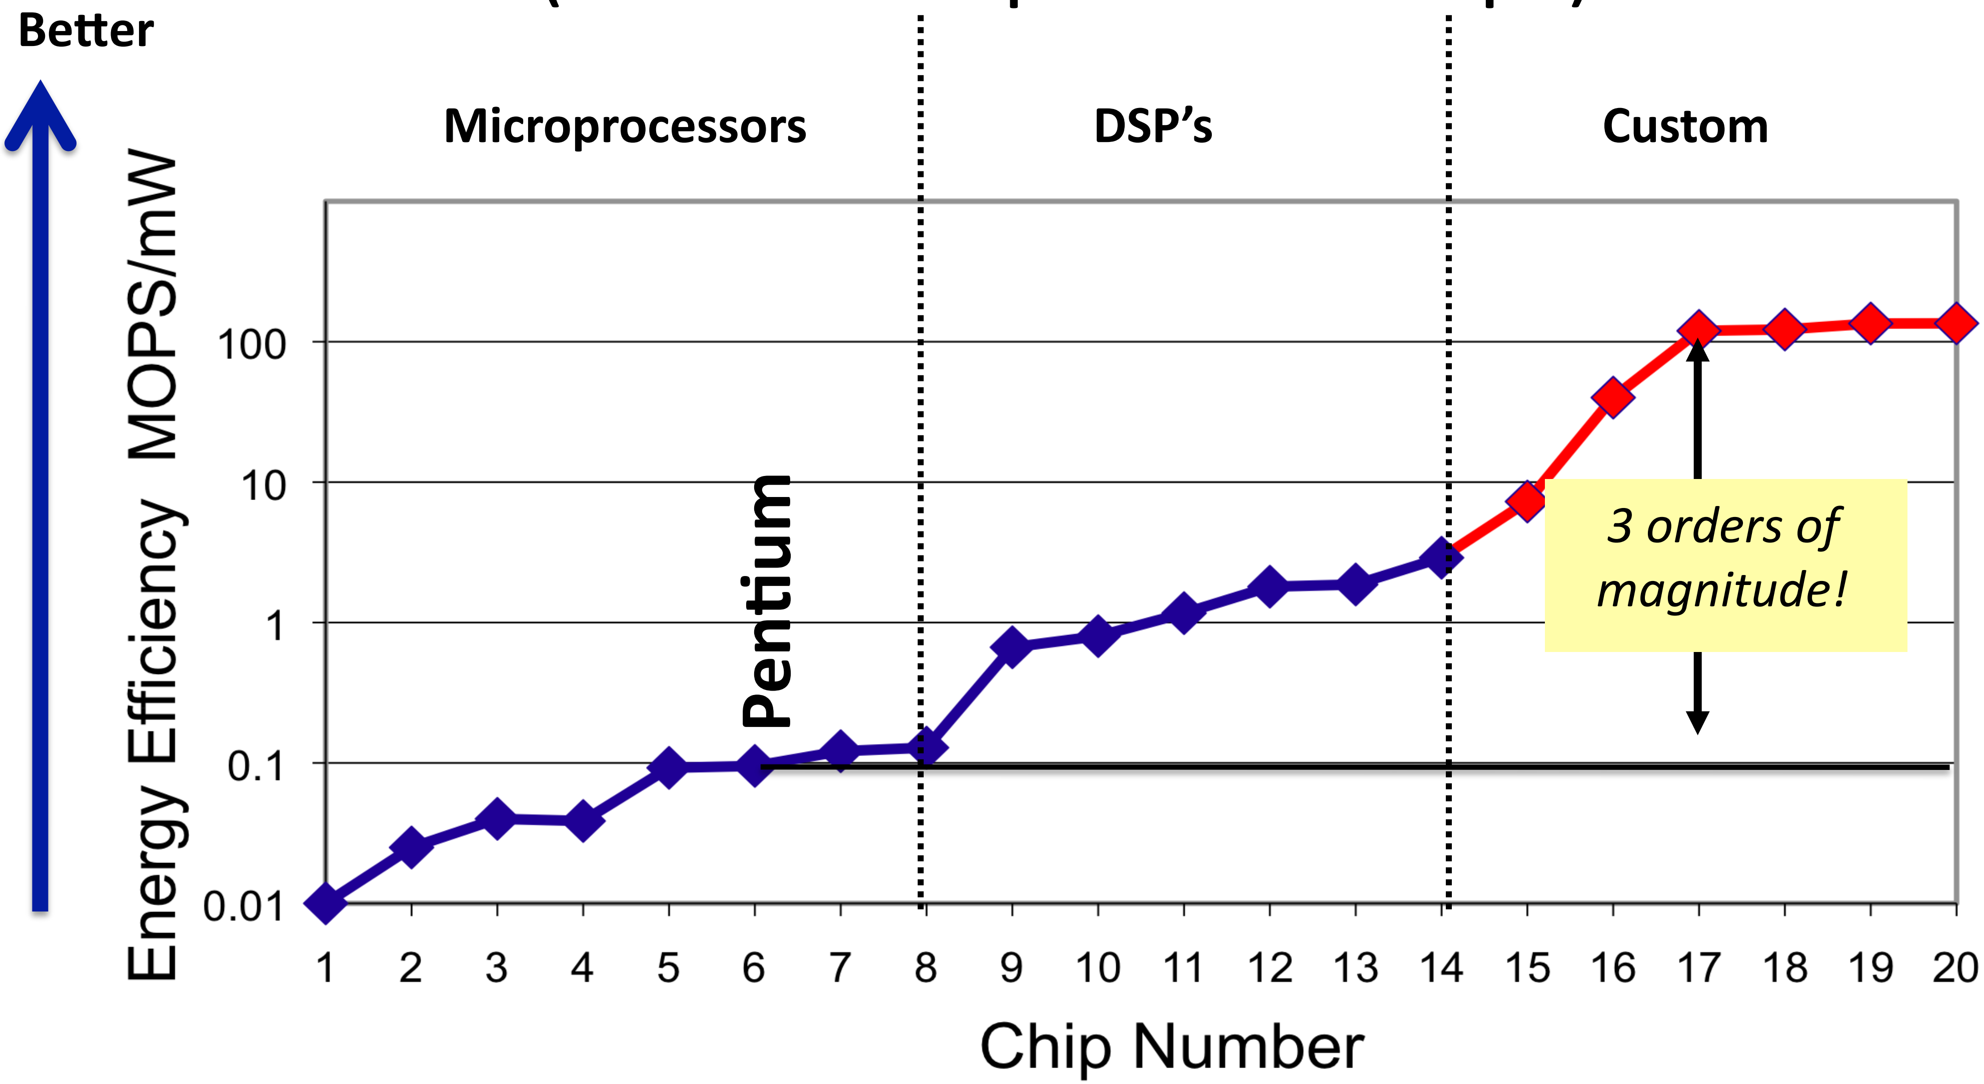

Mihai Budiu, "On the Energy Efficiency of Computation", 2004

Performance ≠ Energy

Energy proportionality

- Avoid idle energy (today 70% of peak in DC's)

Energy minimization

- Perform same task with less energy
  - E.g., specialized hardware
- Exploit slack in energy ➔ Batch tasks/ops

Less energy, same performance

Massive Data meets Energy Wall

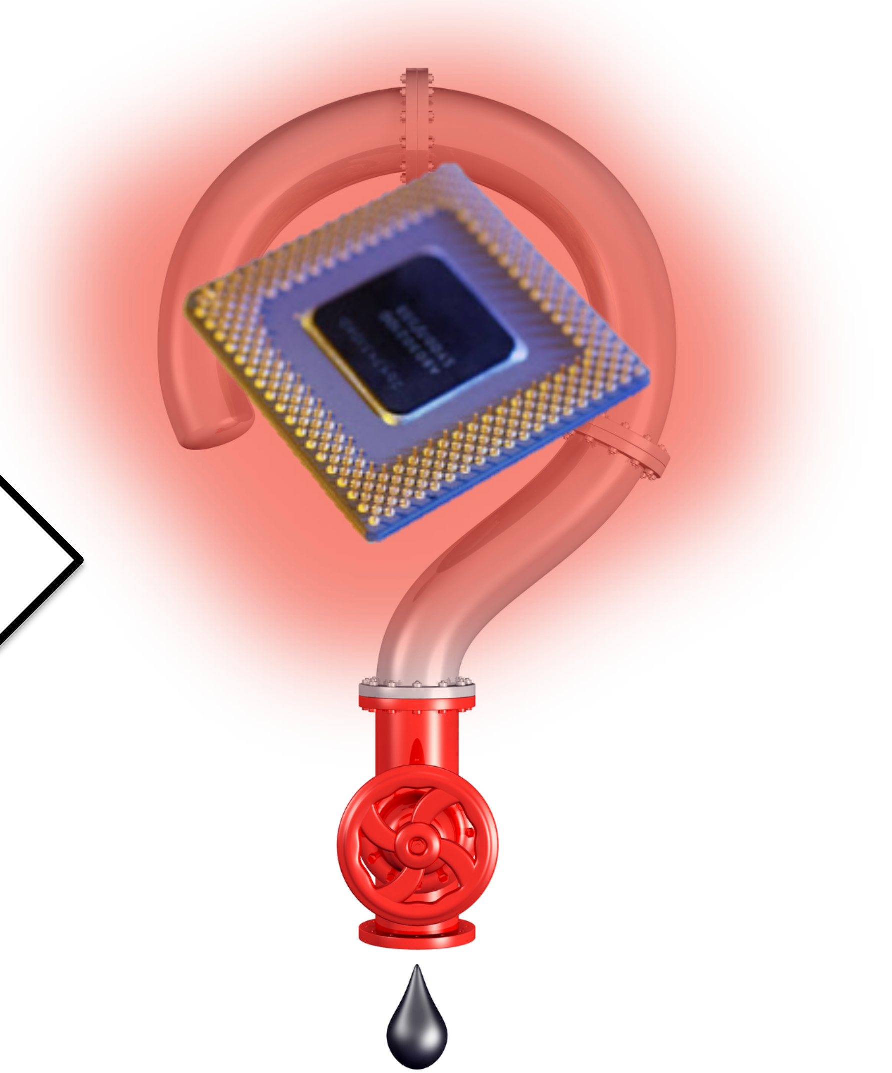

Change in IT's Landscape

- Emergence of Data-Centric Universe
  – IT focus on massive data

- End of "Free Energy"
  – Higher density → higher energy

→Data-Centric Universe meets Energy Wall

What are design implications?

What are the design implications?

Less energy for reliability:

- Trade off HW for less work/energy
- Let it die, detect it, deal with it in SW

Less energy for computation:

- Exploit approximation
- Exploit inherent resilience in data

Temperatures in Datacenters up 30F!

Google, HP, Microsoft

- Pushing temps up
- No perceived change in failure
- High cost reduction

Energy > capital cost:

- Detect failure
- SW fault tolerance takes center stage

Orange - Vendor Spec
Blue- ASHRAE Allowable
Red Outline – New ASHRAE Recommended
Red – Old ASHRAE Recommended

	2004 Version	2008 Version	Most Vendors
Low End Temperature	20°C (68 °F)	18°C (64.4 °F)	10 C
High End Temperature	25°C (77 °F)	27°C (80.6 °F)	35 C
Low End Moisture	40% RH	5.5°C DP (41.9 °F)	20% RH
High End Moisture	55% RH	60% RH & 15°C DP (59 °F DP)	80% RH

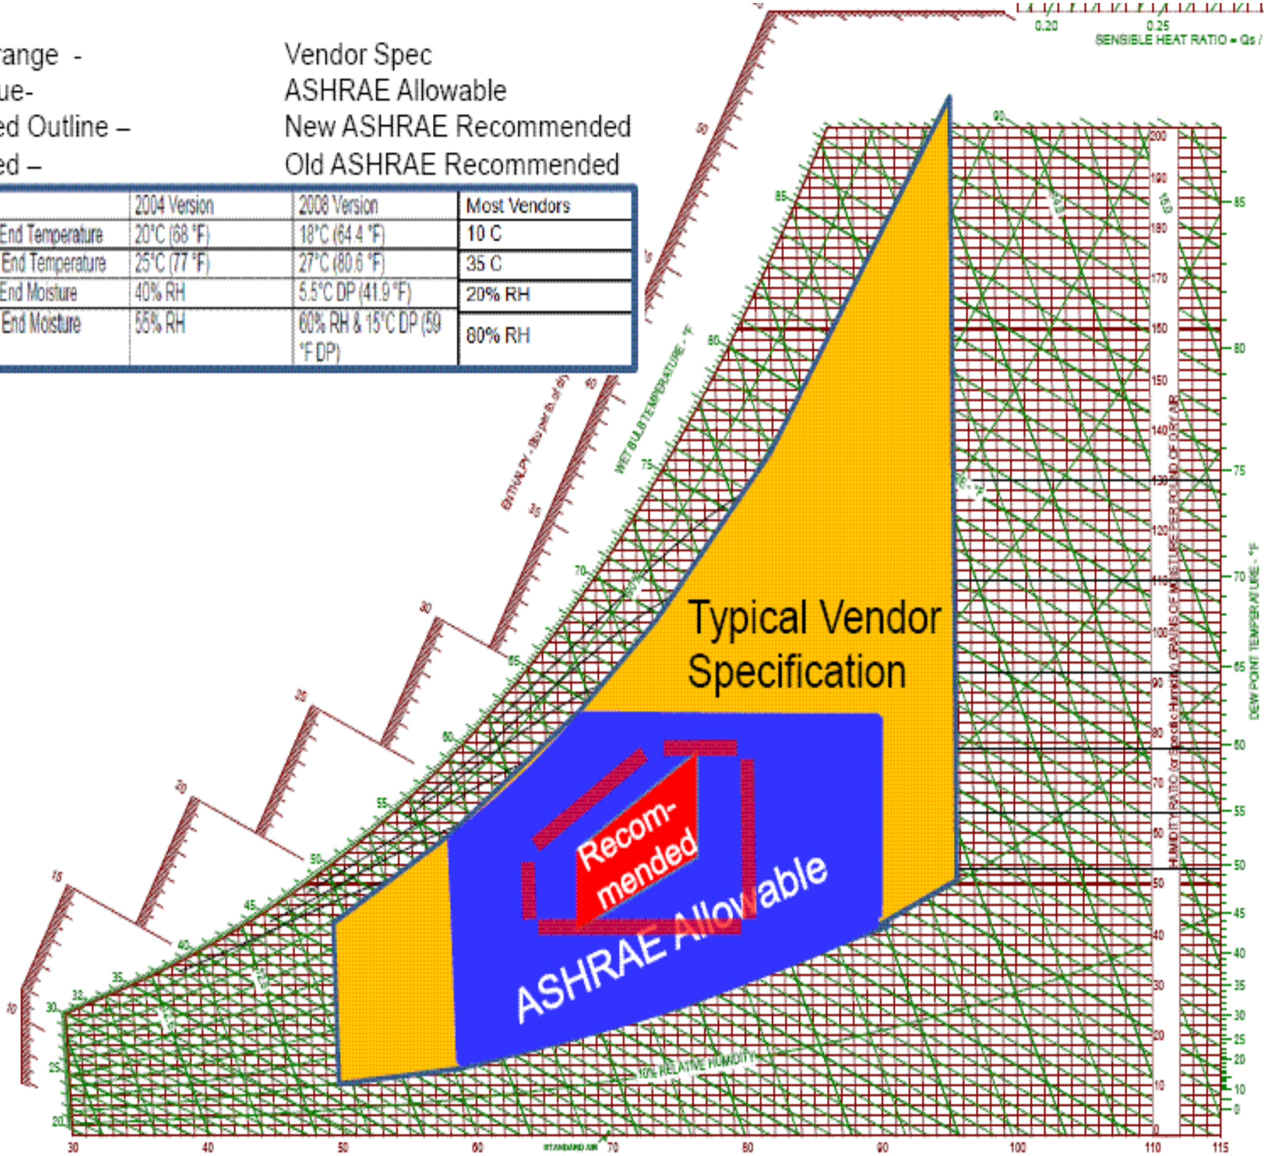

Christian Belady, 2009

Cheaper Reliability

Processors:

- Lightweight detectors
- Fingerprinting, Argus, SWAT revisited for energy
- Sampling DMR

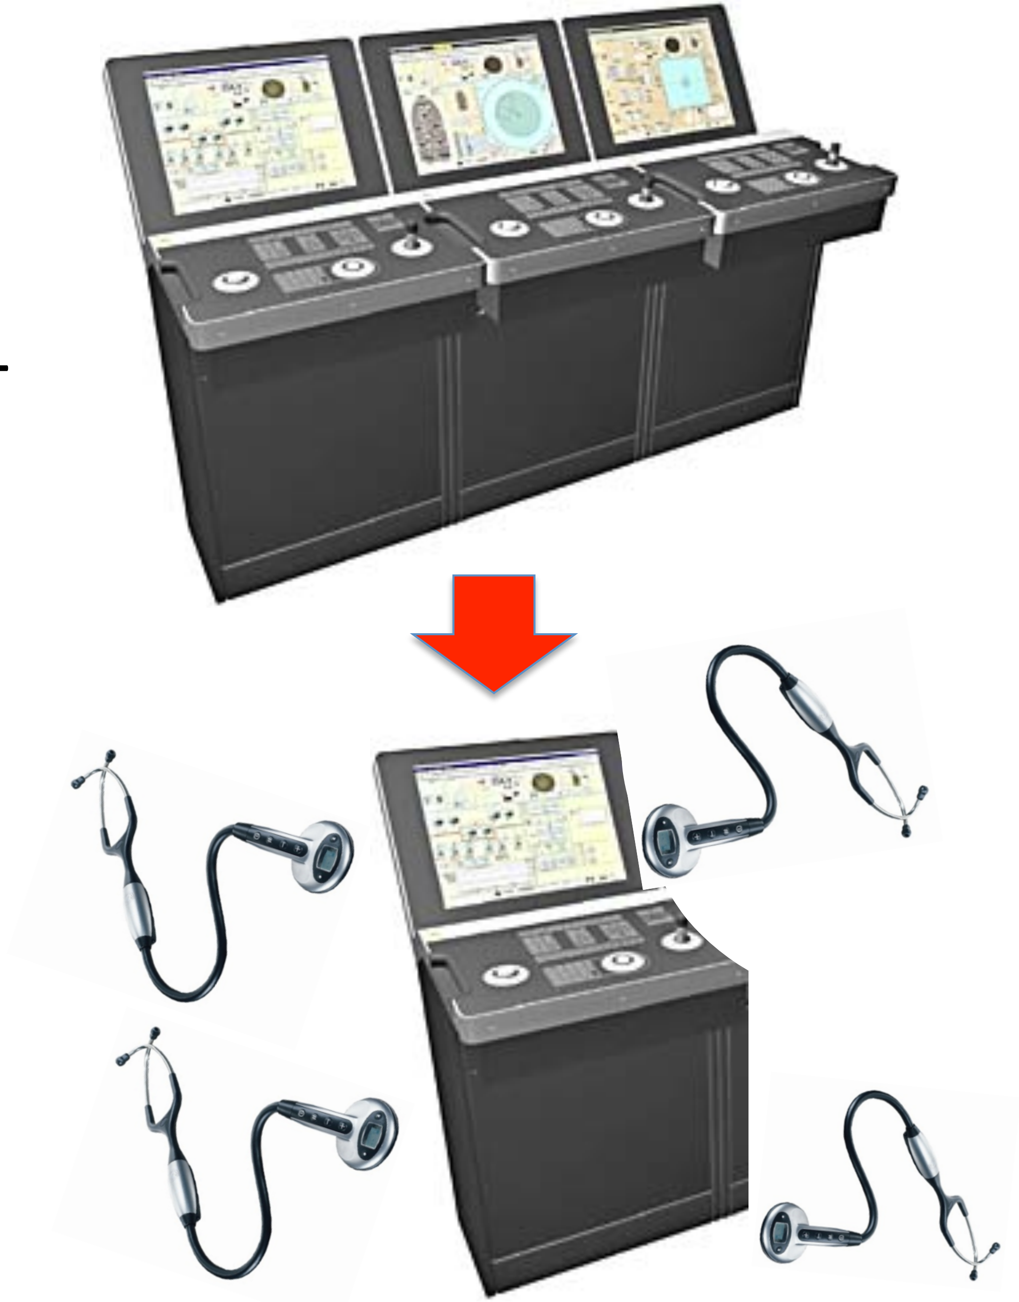

Memory:

- 2D ECC
- Count-based ECC
- Virtualized ECC

Dark Silicon & Reliability

Match made in heaven:

- Provide only necessary coverage
- Specialize reliability (with computation)

Dark silicon:

- Trade off real estate functionality for spare
- Trade off capacity for lower voltages

Exact vs. Probabilistic

Much computation is error-resilient:

- Machine learning/analytics
- Image processing/visual computation
- Audio/speech
- Search

Similarly, two flavors of data

- Exact: affects functionality (pointer address)
- Probabilistic: affects quality (pixels in image)

Perforated (Skipped) Computation

bodytrack benchmark (PARSEC)

- Compiler-driven perforation
- Skip 40% of computation
- Maintains track on head, chest and legs

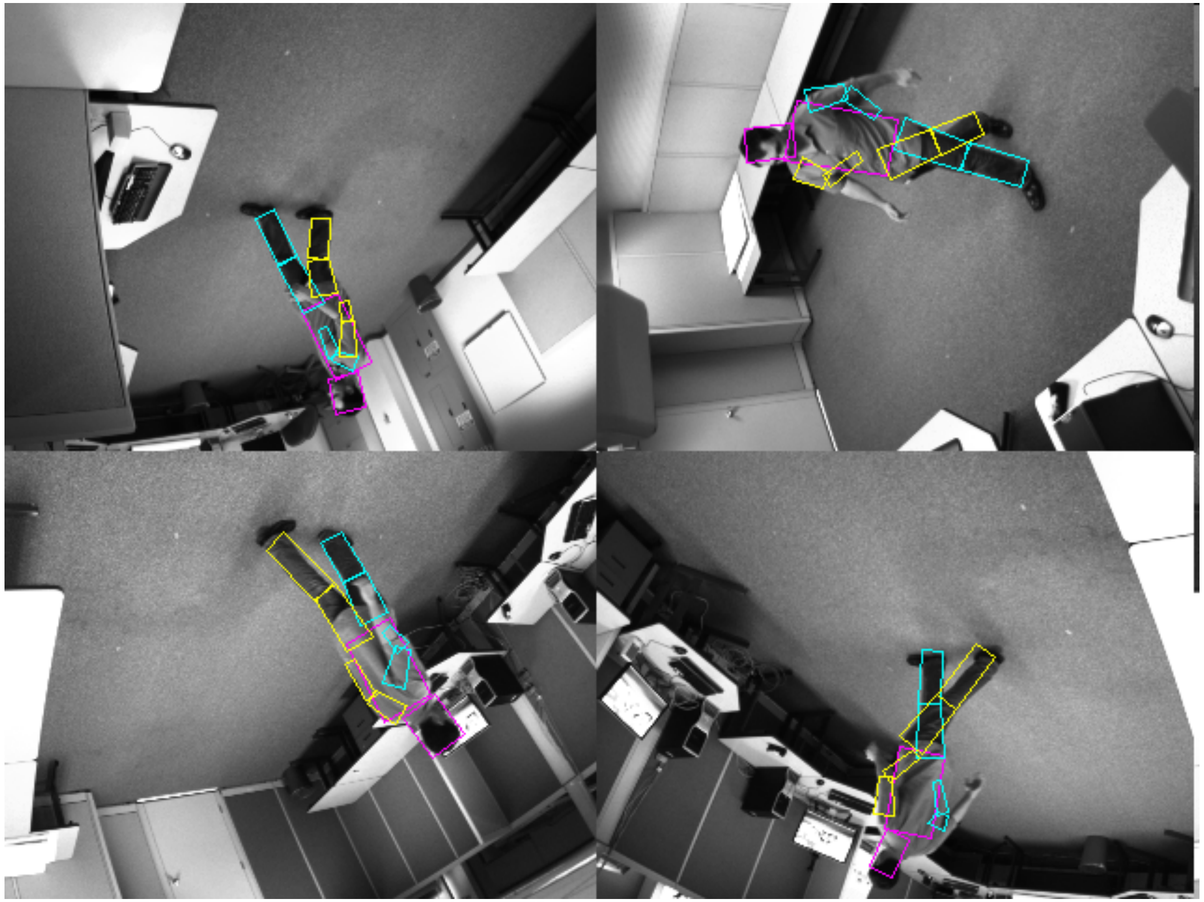

Computation does not have to be exact!

Hoffman et. al., “Using Loop Perforation to Dynamically Adapt Application Behavior to Meet Real-Time Deadlines”, 2010

DeSyRe: Hybrid Reliability

Exploit resilience in computation & data

- Reliable + unreliable substrate
- Partition tasks according to resilience
- Maximize throughput with less energy

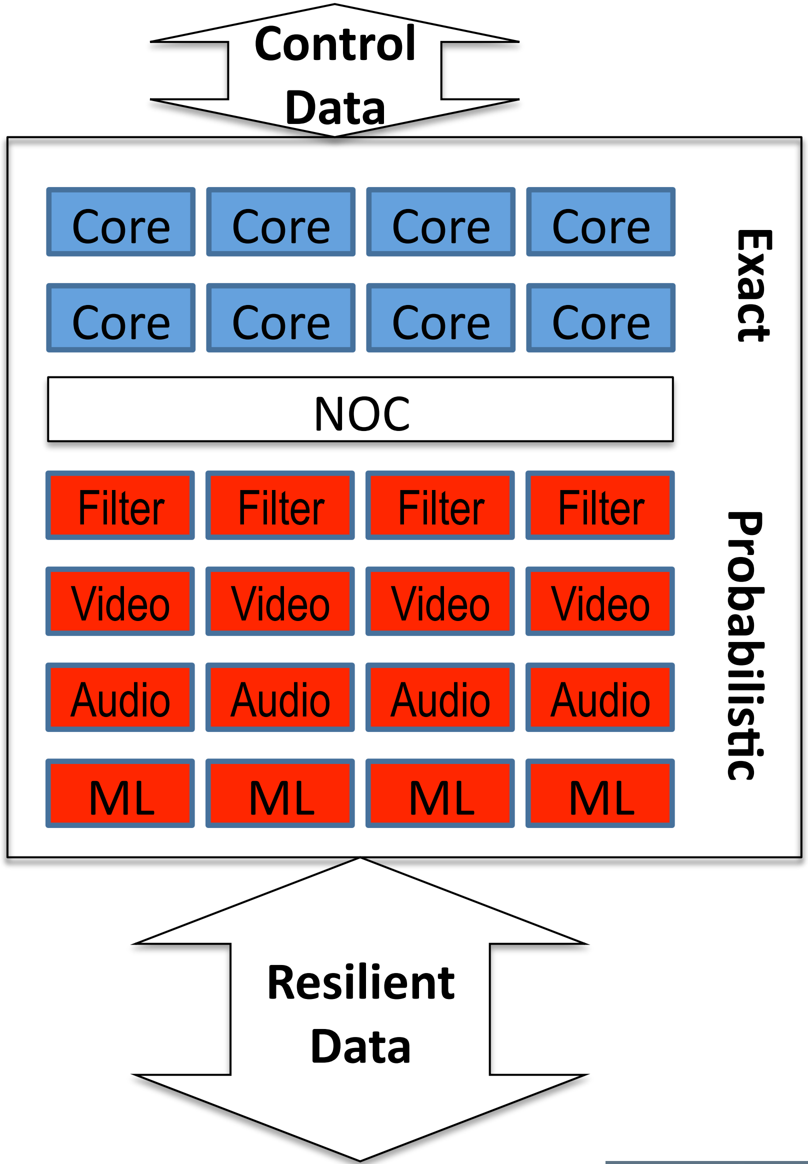

Scaling in Third Dimension

Many technological challenges to overcome

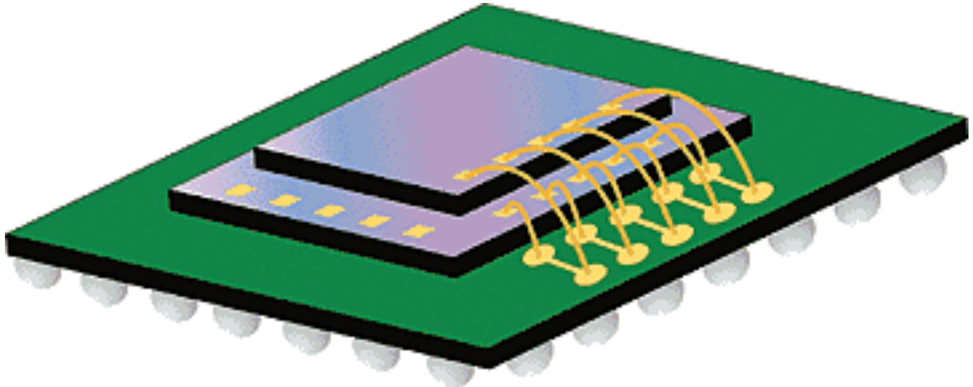

Worse thermal behavior

Great for massive data
- TB/sec bandwidth
- Reduces I/O energy

Scale-Out vs. Scale-Up Workloads

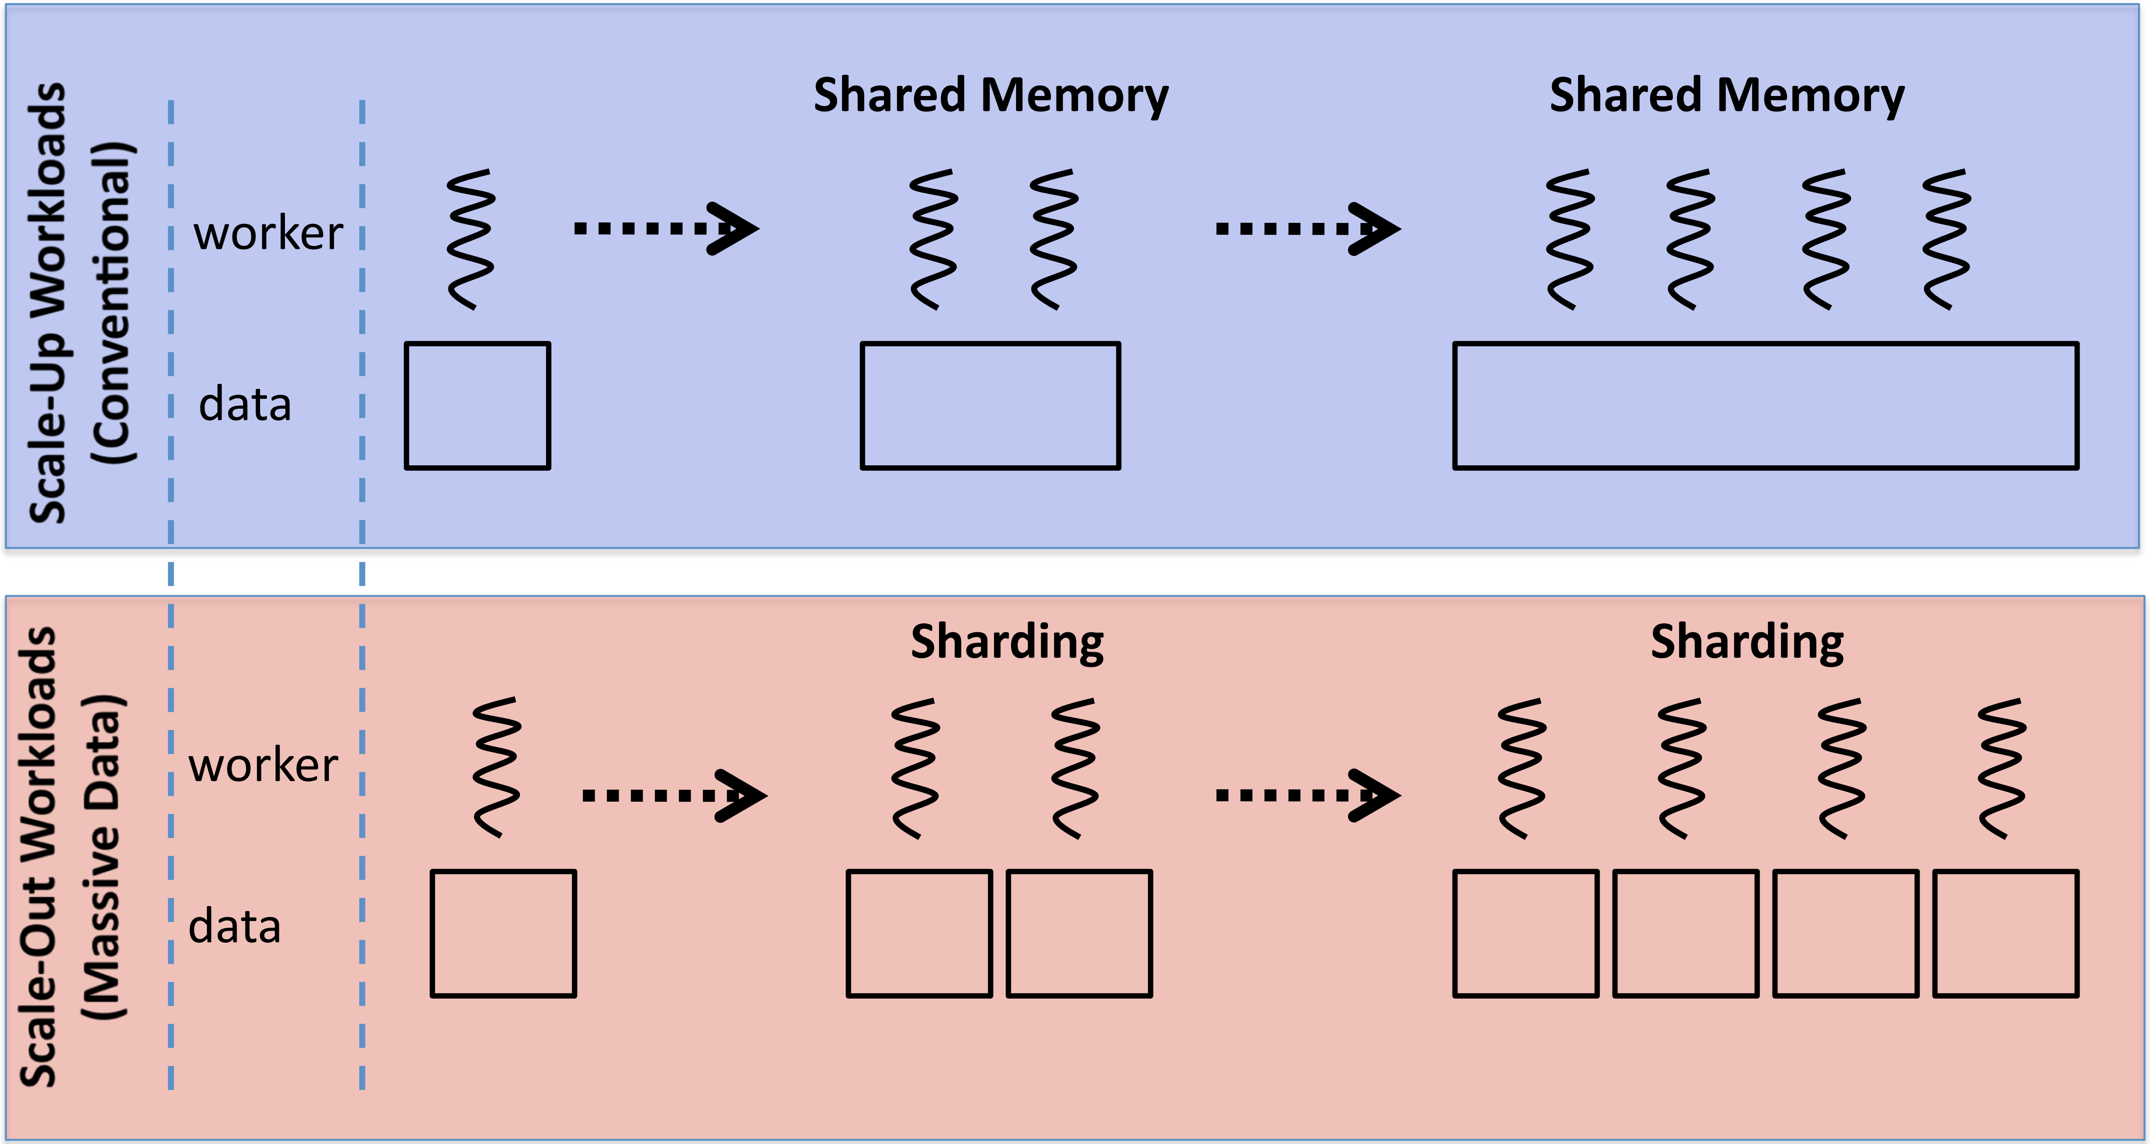

Emerging workloads scale out!

Emerging Workloads are Scale-Out

Examples:

- Data serving (YCSB)
- Streaming
- Search
- Analytics
- Web

Scale-Out vs. Scale-Up Chips

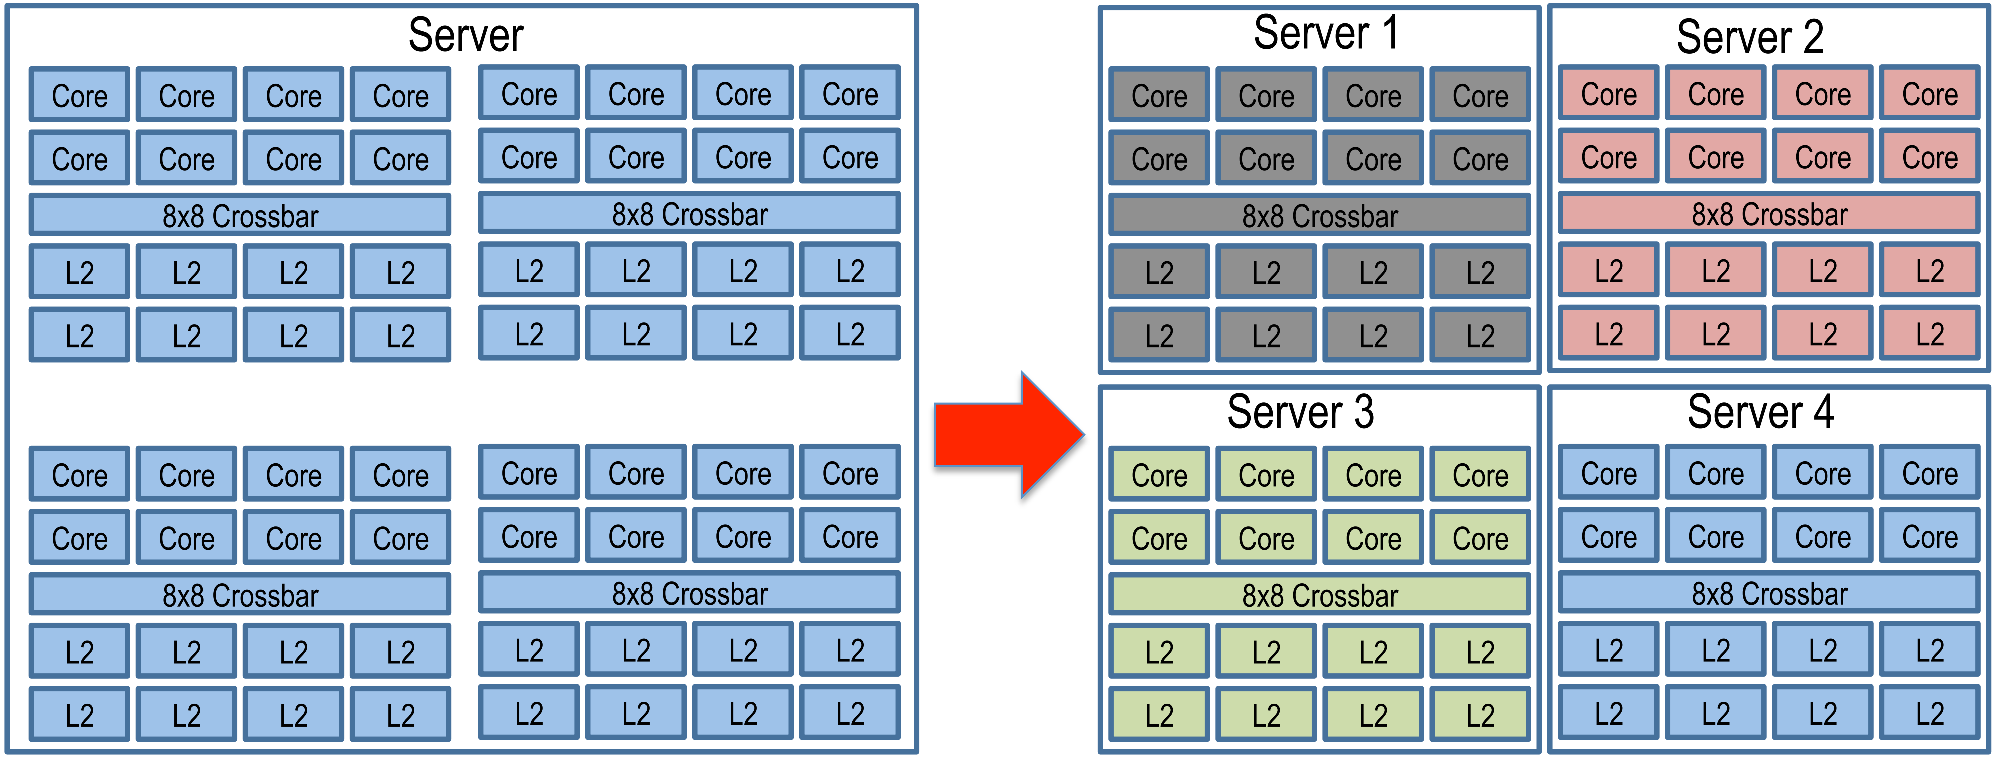

Scale-Up Chip: Conventional Shared Memory **Scale-Out Chip: Clustered Memory**

- Scaling out divides chip among disconnected servers
- Hardware isolation ➔ improved reliability

The EuroCloud Server:
A Scale-Out Chip for Massive Data

(www.eurocloudserver.com)

3D SoC/DRAM:

- 1000x more connectivity
- 10x less system energy
- To run off-the-shelf cloud stack

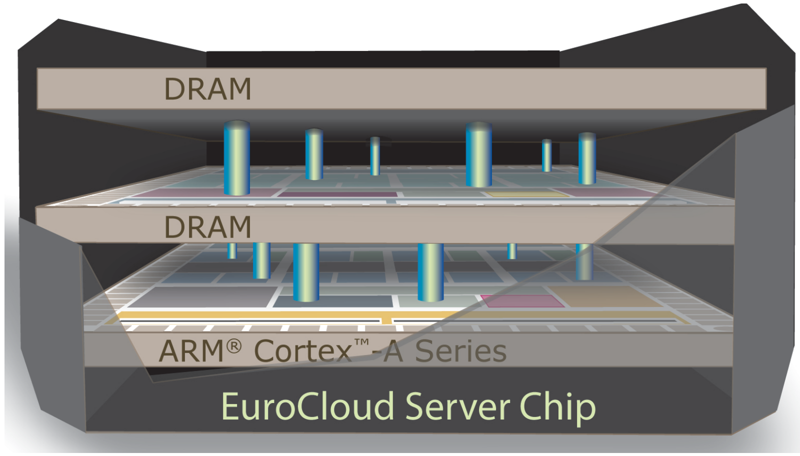

Your Future 1-Watt Datacenter Chip

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

ARM

NOKIA Connecting People

imec

UCyprus

Bringing it All Together

- IT is changing everything & itself changing
- Future: Plow massive data with minimal energy
- Our way of computing is inefficient
- Are our systems too exact, too robust?

Oh, I can be reliable and for an extra 300 hundred a week can be efficient too.

Thank You!

For more information please visit us at

ecocloud.ch

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE