The Future of Advanced Computing: It's Uncertain

Dan Reed
Vice President for Research and Economic Development
University Computational Science and Bioinformatics Chair
Computer Science, Electrical Engineering & Computer Engineering, and Medicine

THE UNIVERSITY OF IOWA

dan-reed@uiowa.edu www.hpcdan.org

Discussion roadmap

Changing nature of science

Technology lessons and exponentials

Re-conceptualizing advanced computing

Thoughts on the future

Human Genome Functional Elements

Quantum mechanics spacetime gravity

$$W = \int_{k<\Lambda} [Dg][DA][D\psi][D\Phi] \exp\left\{ i \int d^4x\sqrt{-g}\left[\frac{m_p^2}{2} R \right.\right.$$

$$\left.\left. -\frac{1}{4} F^a_{\mu\nu} F^{a\mu\nu} + i\bar{\psi}^i \gamma^\mu D_\mu \psi^i + \left(\bar{\psi}^i_L V_{ij} \Phi \psi^j_R + \text{h.c.} \right) - |D_\mu \Phi|^2 - V(\Phi) \right] \right\}$$

Other forces matter Higgs

http://www.preposterousuniverse.com

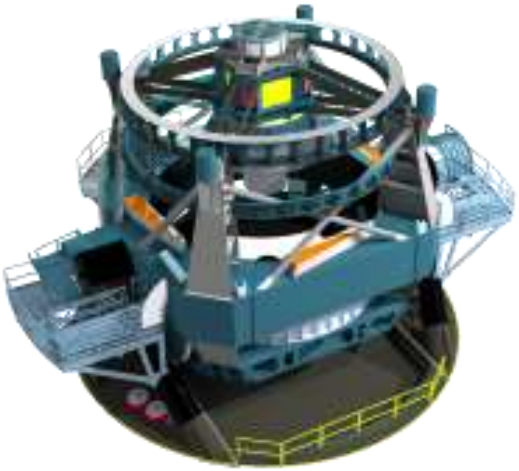

Large Synoptic Survey Telescope

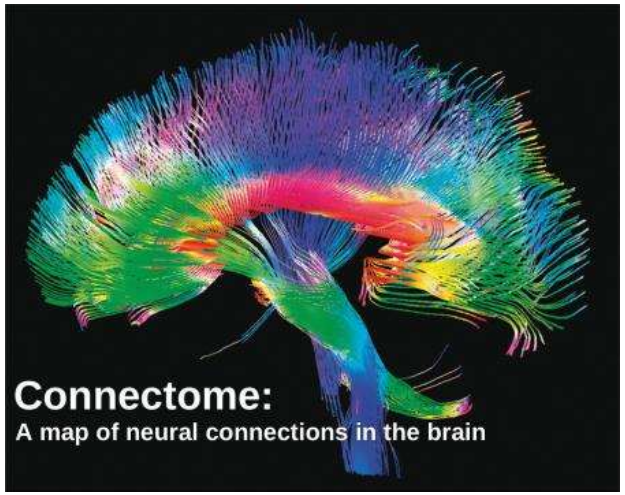

Brain Connectome

Ocean Observing System

Square Kilometer Array

Square Kilometer Array (SKA)

BIG, Big, big data

Structure

- SKA-MID (South Africa) with ~2000 15m dishes
- SKA LOW (Australia) with ~1 million dipole antennas
- Spread over 3000 km with construction in 2018 …
- 50 MHz – 14 GHz coverage; 100X sensitivity; 10^6 faster

Science

- Galaxy evolution and dark energy
- Gravitational waves
- Dark Ages look back

Data and processing

- From 100 PB/day to 10,000 PB/day
- Real-time reduction at exascale levels
  - World's biggest FFTs

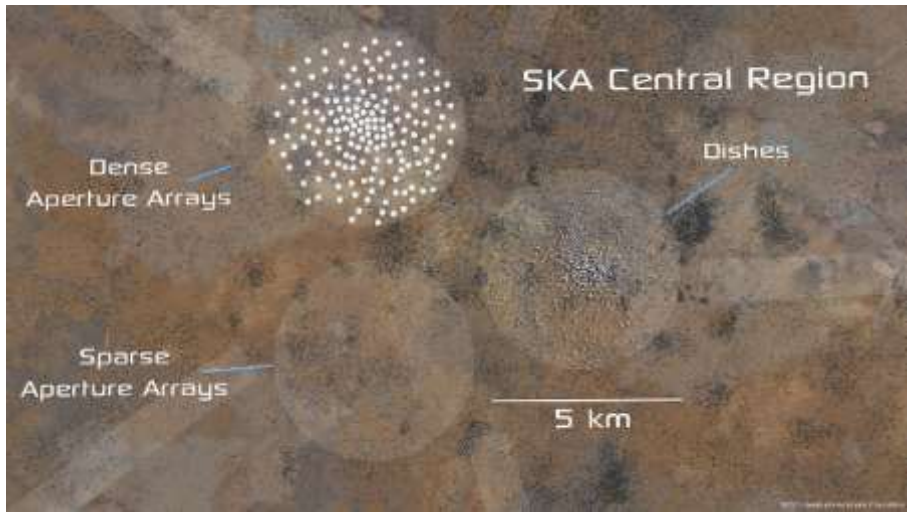

Ocean Observing Initiative (OOI)
Heterogeneous data

Structure
- Coastal and Global Scale Nodes (CGSNs)
  - Buoys, moorings, cables and AUVs
- Regional Scale Nodes (RSNs)
  - 900 km of cabled sensor arrays

Science
- Ocean-atmosphere exchange
- Climate variability
- Turbulent mixing and biophysical interactions
- Coastal ocean dynamics
- Geophysics and geodynamics

Data and processing
- Extremely heterogeneous data
- Complex software and data products

Computing transitions: Riding the economic wave

Mainframes | **Vectors** (1980s) | **MPPs** (1990s) | **Clusters & Grids** (2000s) | Clouds, Big Data and Devices

Google

Windows Azure

amazon web services™

Relational DBMS | **Object-oriented DBMS** | **"No SQL" DBMS** | **Map-Reduce**

NVIDIA

intel

The talent in data analytics have shifted from science to companies. We can't compete.

Astronomy researcher

The University of Iowa

Quantitative begets qualitative

Seagate 2 TB Disk

200 MB disk pack (Wikipedia)

Four computing capabilities

- Networking – delivering questions and answers
- Computation – transforming data into insights
- Data access – timely access to data by computations
- Storage – long-term data retention

Three defining attributes …

- Speed, capacity and cost

… and their ratios determine viability

Exponentials are very deceptive

- Early, the future looks like the past
- *Suddenly, everything is different*

Generation after generation
Disrupted from below by 10X

Volume
Unit price
Market size

Mainframes
IBM S/360

Minicomputers
VAX 11/780

Workstations
SUN 3/50

Personal Computers
IBM PC

Windows PCs

Tablets/Smartphones
iPad/iPhone/Galaxy

Frozen here?

Market disruption
Performance/$
Societal impact

Chip feature sizes

No exponential is forever (except in the textbooks)

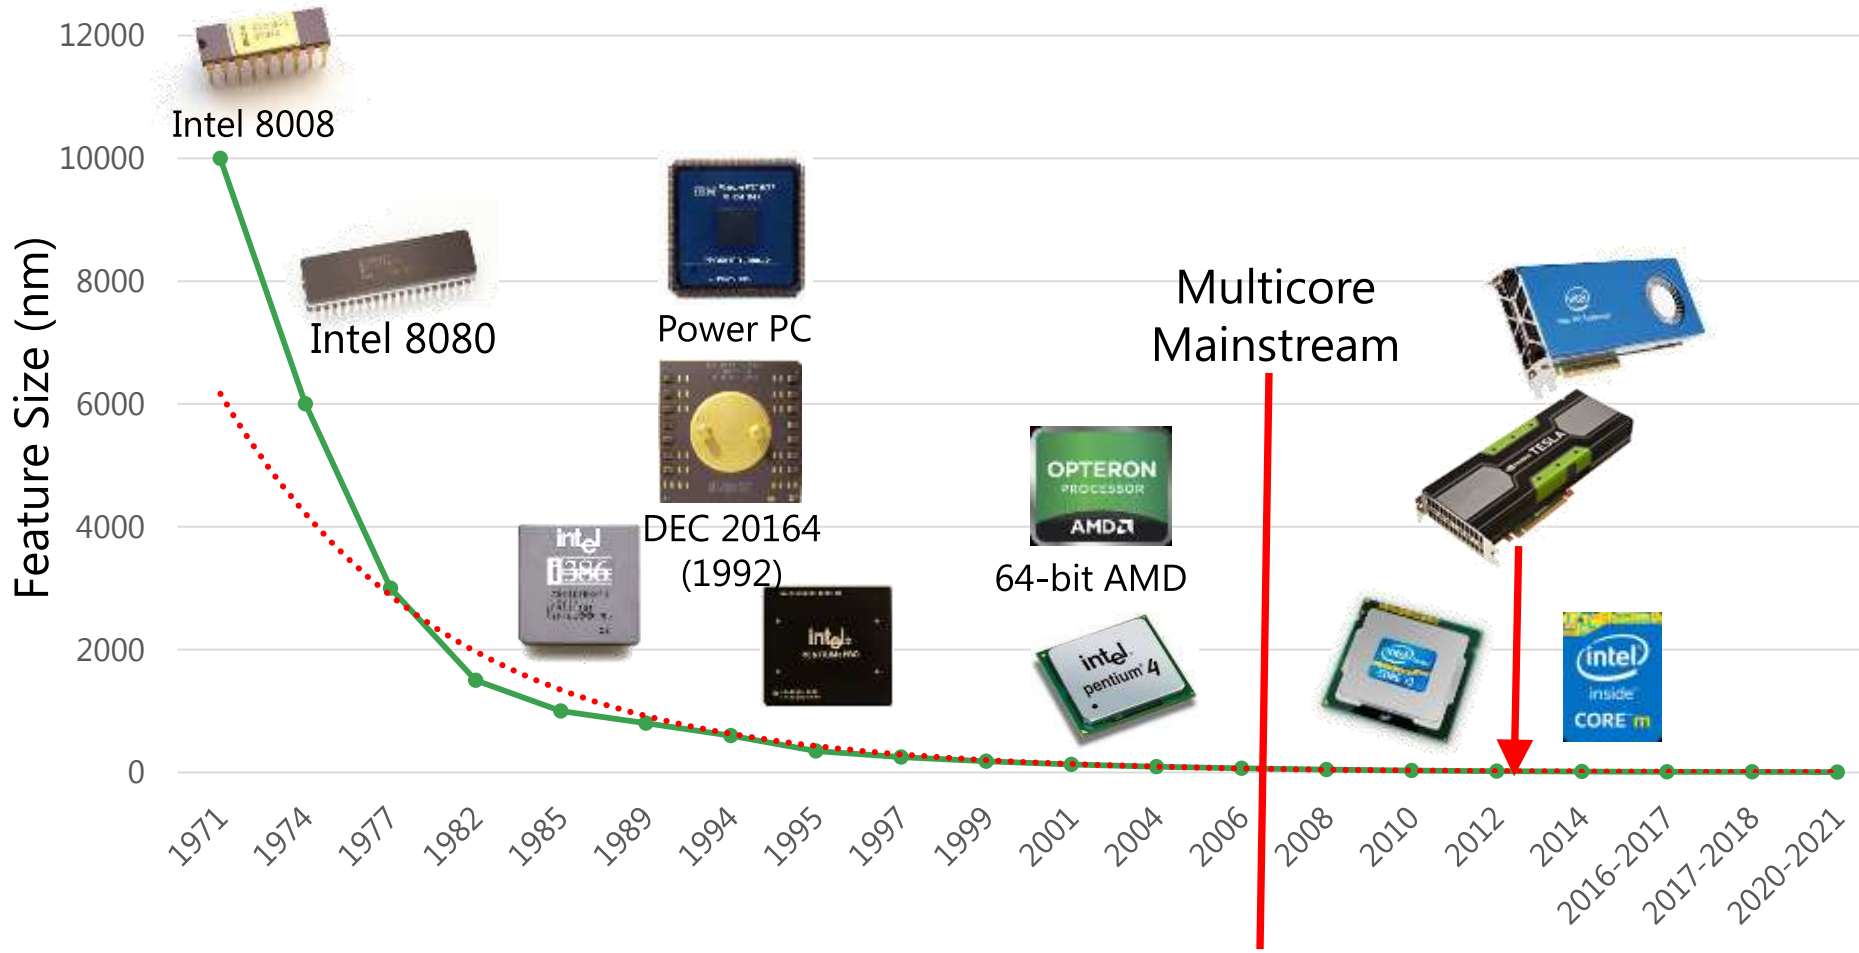

Room at the bottom: less and less

Yeah, we did that – Dennard scaling has come and gone

Why can't we make them [computers] very small, make them of little wires, little elements – and by little, I mean *little*. For instance, the wires should be 10 or 100 atoms in diameter, and the circuits should be a few thousand angstroms across. Everybody who has analyzed the logical theory of computers has come to the conclusion that the possibilities of computers are very interesting – if they could be made to be more complicated by several orders of magnitude.

Richard Feynman
December 29, 1959

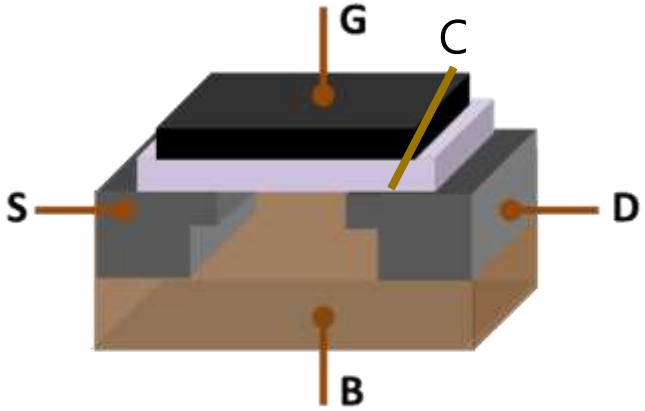

MOSFET (Wikipedia)
G (gate), S (source)
D (drain) and C (channel)

$$CV^2F + S$$

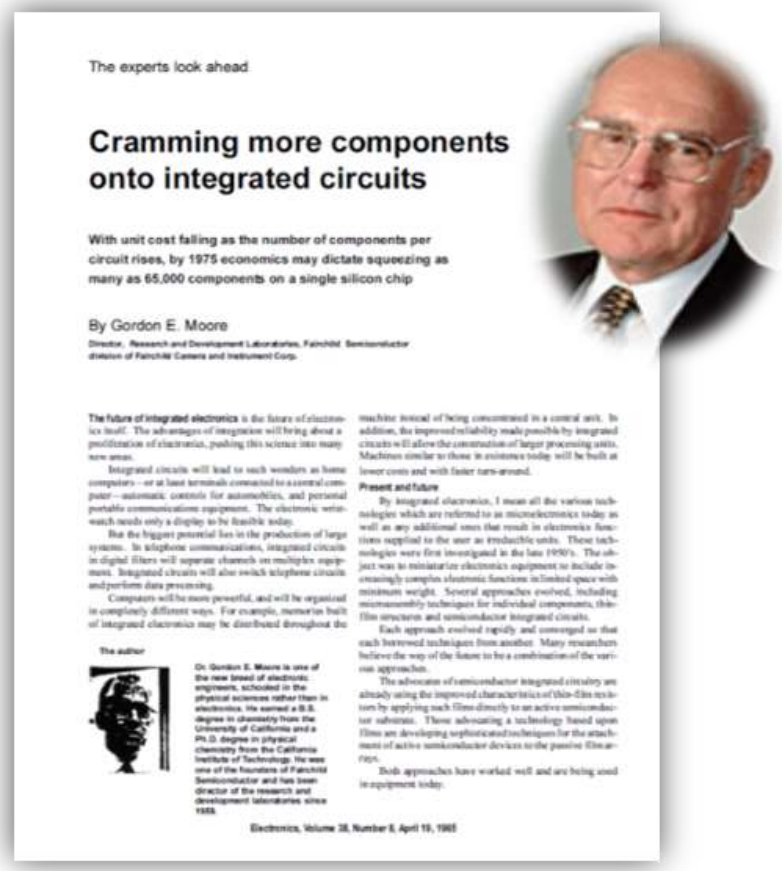
The experts look ahead

Cramming more components onto integrated circuits

With unit cost falling as the number of components per circuit rises, by 1975 economics may dictate squeezing as many as 65,000 components on a single silicon chip

By Gordon E. Moore

Director, Research and Development Laboratories, Fairchild Semiconductor division of Fairchild Camera and Instrument Corp.

The future of integrated electronics is the future of electronics itself. The advantages of integration will bring about a proliferation of electronics, pushing this science into many new areas.

Integrated circuits will lead to such wonders as home computers—or at least terminals connected to a central computer—automatic controls for automobiles, and personal portable communications equipment. The electronic wristwatch needs only a display to be feasible today.

But the biggest potential lies in the production of large systems. In telephone communications, integrated circuits in digital filters will separate channels on multiplex equipment. Integrated circuits will also switch telephone circuits and perform data processing.

Computers will be more powerful, and will be organized in completely different ways. For example, memories built of integrated electronics may be distributed throughout the machine instead of being concentrated in a central unit. In addition, the improved reliability made possible by integrated circuits will allow the construction of larger processing units. Machines similar to those in existence today will be built at lower costs and with faster turn-around.

Present and future

By integrated electronics, I mean all the various technologies which are referred to as microelectronics today as well as any additional ones that result in electronics functions supplied to the user as irreducible units. These technologies were first investigated in the late 1950's. The object was to miniaturize electronics equipment to include increasingly complex electronic functions in limited space with minimum weight. Several approaches evolved, including microassembly techniques for individual components, thin-film structures and semiconductor integrated circuits.

Each approach evolved rapidly and converged so that each borrowed techniques from another. Many researchers believe the way of the future to be a combination of the various approaches.

The advocates of semiconductor integrated circuitry are already using the improved characteristics of thin-film resistors by applying such films directly to an active semiconductor substrate. Those advocating a technology based upon films are developing sophisticated techniques for the attachment of active semiconductor devices to the passive film arrays.

Both approaches have worked well and are being used in equipment today.

The author

Dr. Gordon E. Moore is one of the new breed of electronic engineers, schooled in the physical sciences rather than in electronics. He earned a B.S. degree in chemistry from the University of California and a Ph.D. degree in physical chemistry from the California Institute of Technology. He was one of the founders of Fairchild Semiconductor and has been director of the research and development laboratories since 1959.

Electronics, Volume 38, Number 8, April 19, 1965

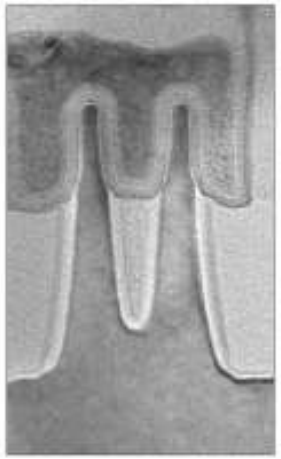

14 nm 2nd Generation Tri-gate Transistor

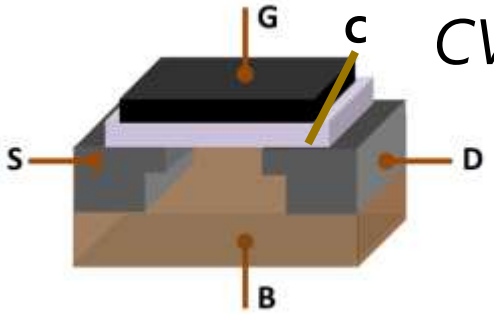

$CV^2F + S$

Cryofab

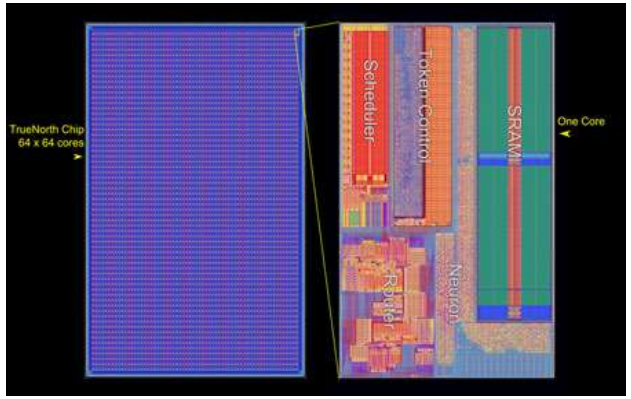

IBM True North chip
4096 cores x 256 neurons

Inflection points nearby

A bit of inflection point history
NCSA terascale and PlayStation2 clusters (~2001)

1 TF IA-32 Pentium III cluster (Platinum)
- 512 1 GHz dual processor nodes
- Myrinet 2000 interconnect and 5 TB of RAID storage
- 30th on Top500 list (world's fastest Linux cluster)

1 TF IA-64 Itanium cluster (Titan)
- 164 800 MHz dual processor nodes
- Myrinet 2000 interconnect

IBM

intel

Myricom

NCSA

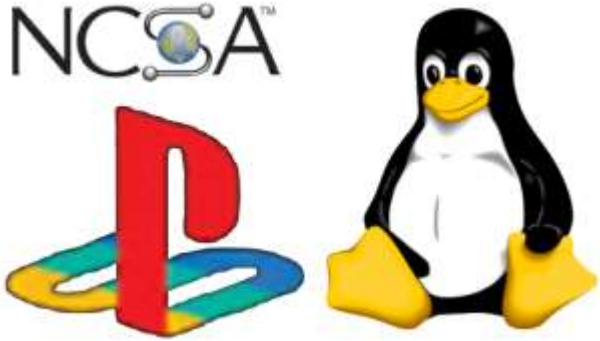

THE UNIVERSITY OF IOWA

A lesson from another domain

Airspeed records (kilometers/hour)

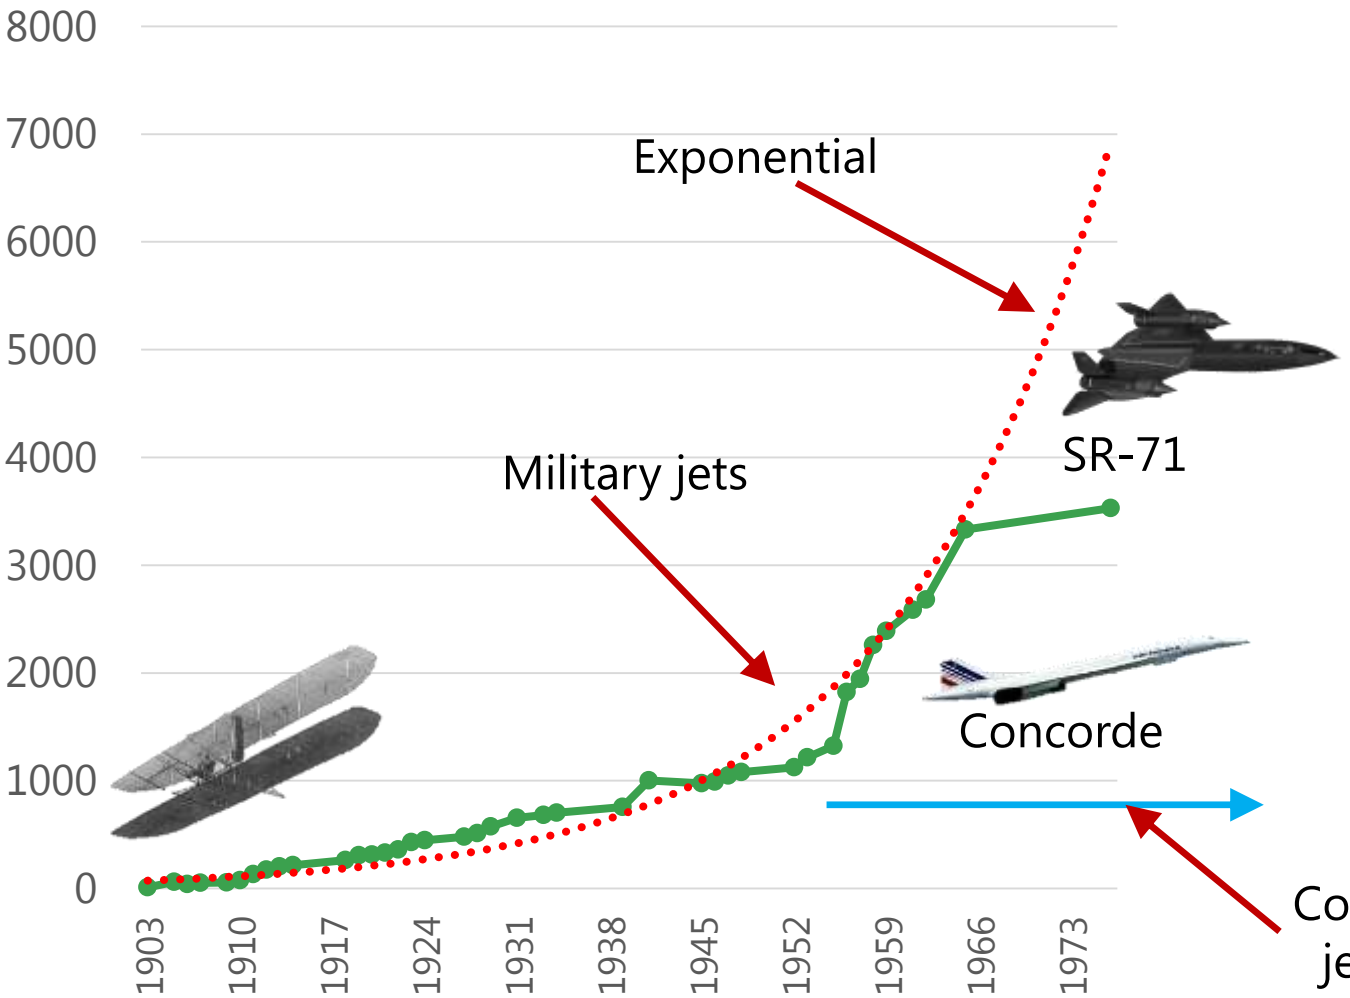

Loci of airplane innovation shifted

- Capacity
- Range
- Safety
- Yield management
- Avionics
- Fuel efficiency
- Carbon composites

… but it has never stopped

Computing loci are also shifting

Diverging cultures

Technical application complexity is rising

- Multidisciplinary fusion
- Temporal and spatial adaptation
- Data assimilation and processing

… along with multiple optimization axes

- Massive parallelism with heterogeneous cores
- Resilience/reliability at large scale
- Energy optimization for utility

C, Fortran, C++, MPI, OpenMP

Python, Ruby, R

Cloud/Web Services

Technical and mainstream software development have diverged

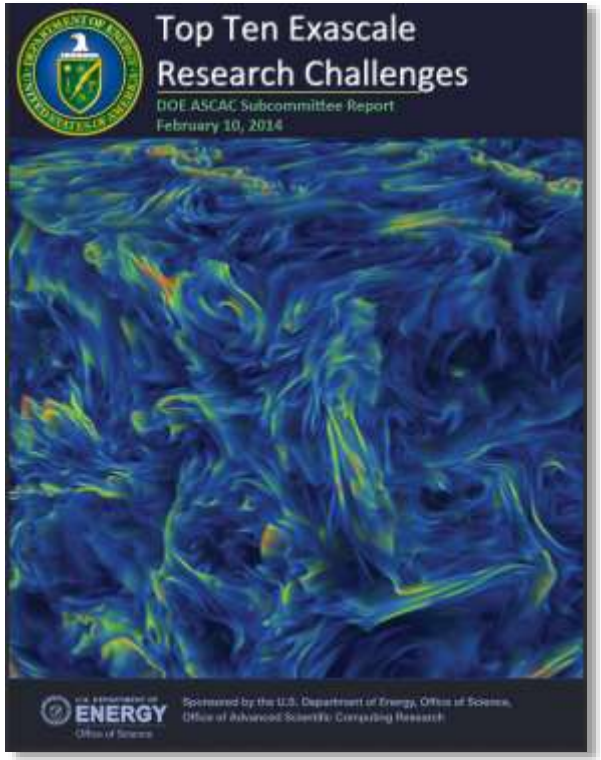

Divergent ecosystems

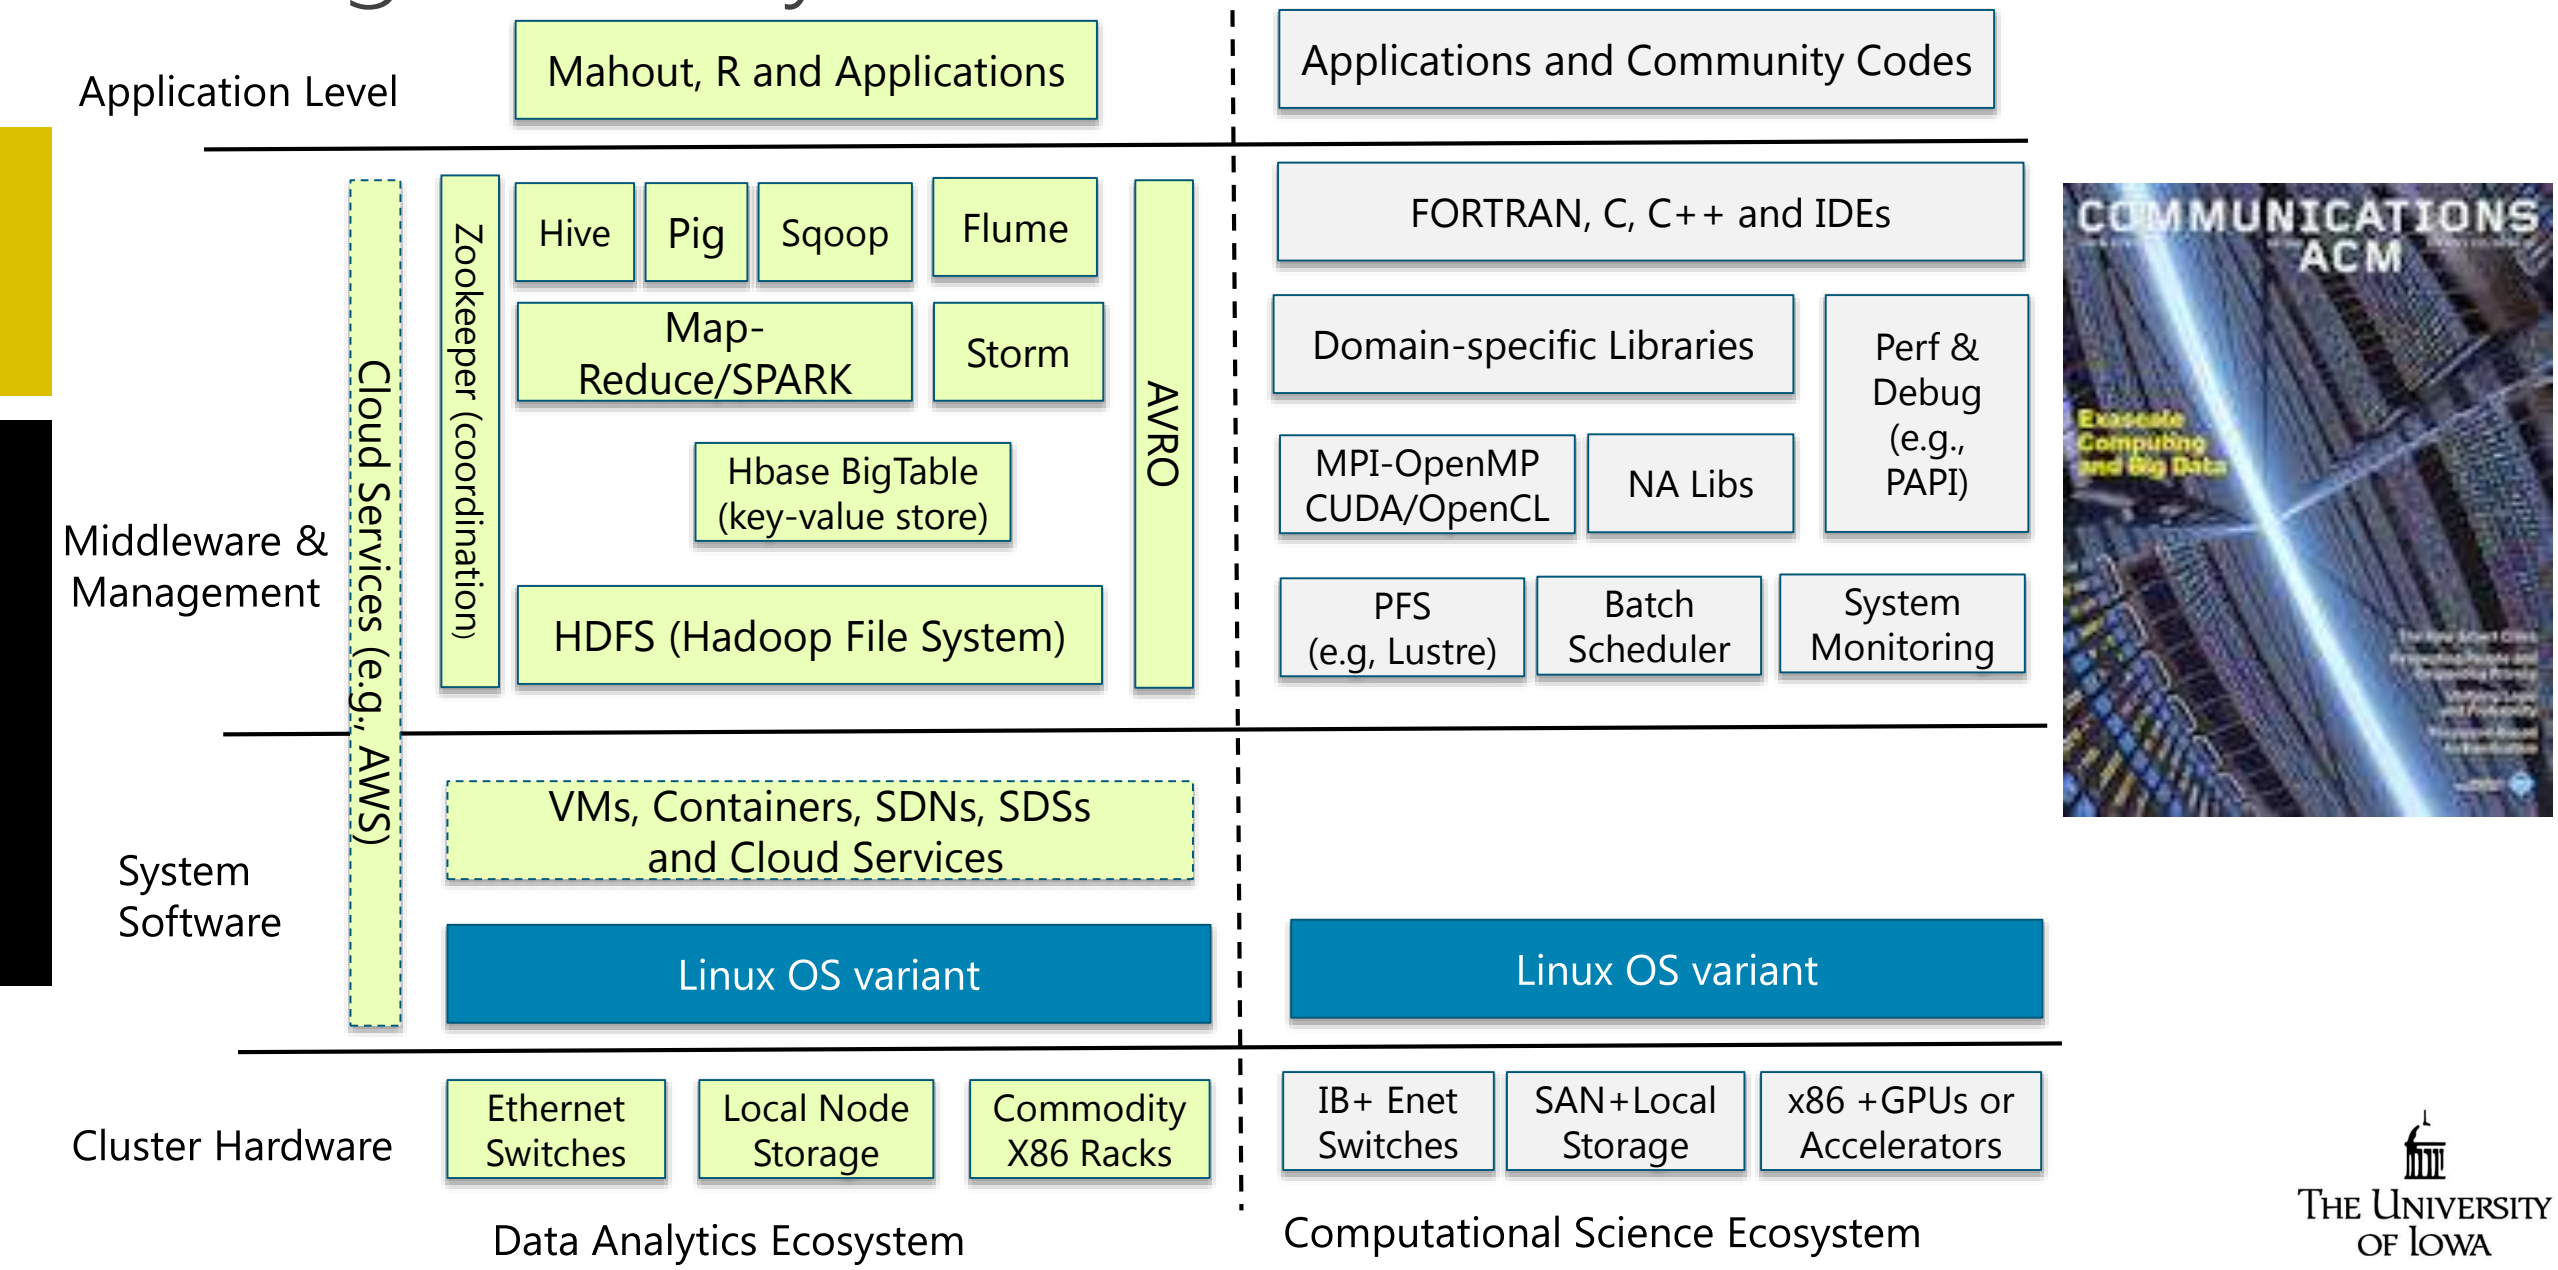

Recommender techniques

Item hierarchy (Amazon)

- You bought a Kindle™, you'll want a cover

Attributes (Pandora)

- You like 70's pop, you'll like *Simon and Garfunkel*

Item similarity (Netflix)

- You liked *Batman*, you'll like *Spiderman*

User similarity (Walmart)

- People who buy beer generally buy chips

Social network (Linkedin)

- Your connections liked this job posting, so will you

Model based (HPC challenges and needs)

- ***Training, singular value decomposition (SVD), support vector machines (SVM)***
- ***Supervised and unsupervised learning …***

A machine learning revolution is underway – if we let it

Big data and massive parallelism

- Supervised machine learning – inference from labeled training sets
- Unsupervised machine learning – finding structure in data without labels

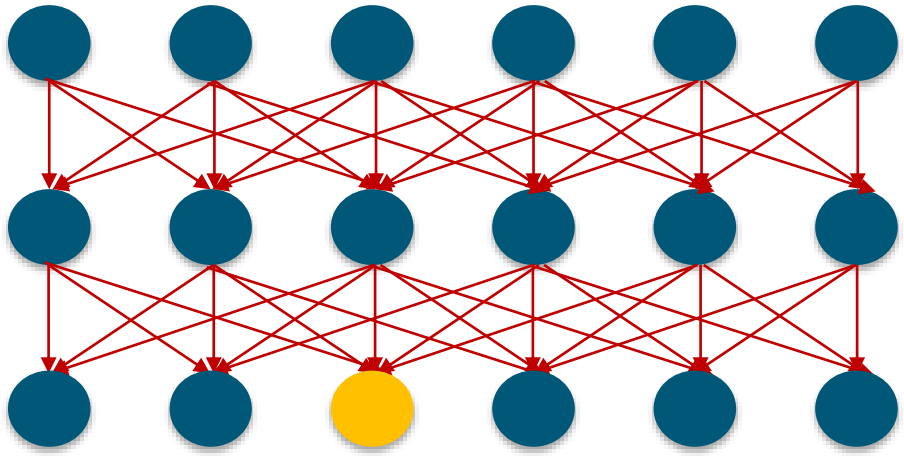

Spiral Galaxy

Labels

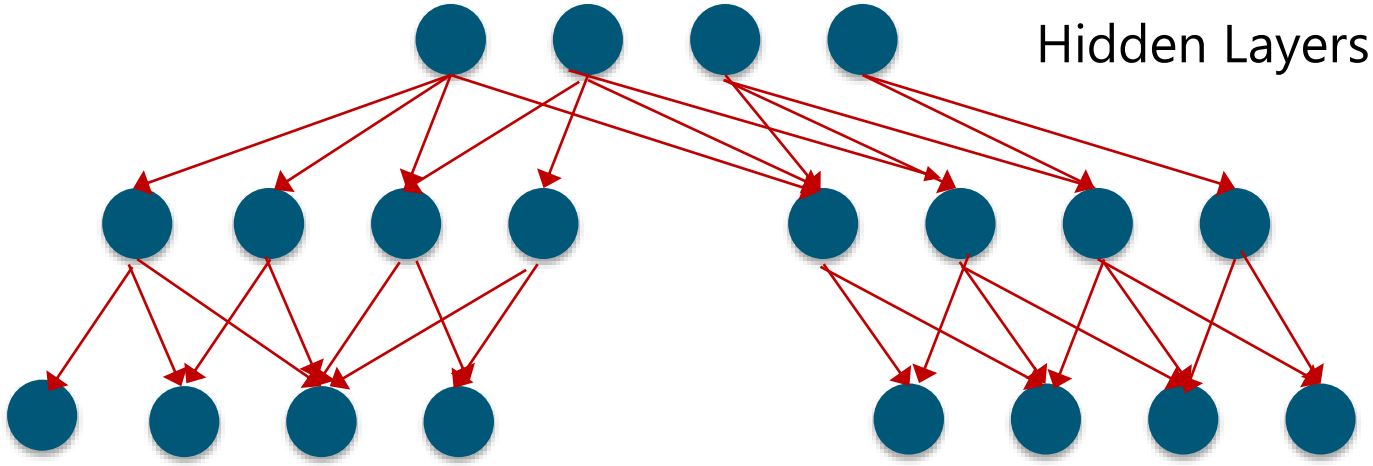

Detected Features

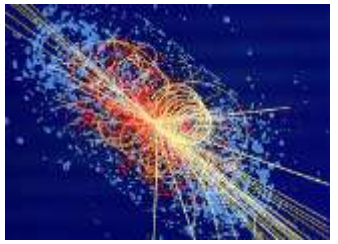

Input Data

Source: Dennis Gannon

Have you ever …

Requested compute and storage for *ten years?*

Logged onto a node and killed processes just to see what would happen?

Wished you could load containers rather than just applications?

Found your code performance limited by the I/O bandwidth of a Raspberry Pi?

Thought SAN was just a typo in a message meant for Sam?

Asked your system for recommendations?

Wondered why R came after S and C doesn't matter?

Follow the money and the users …

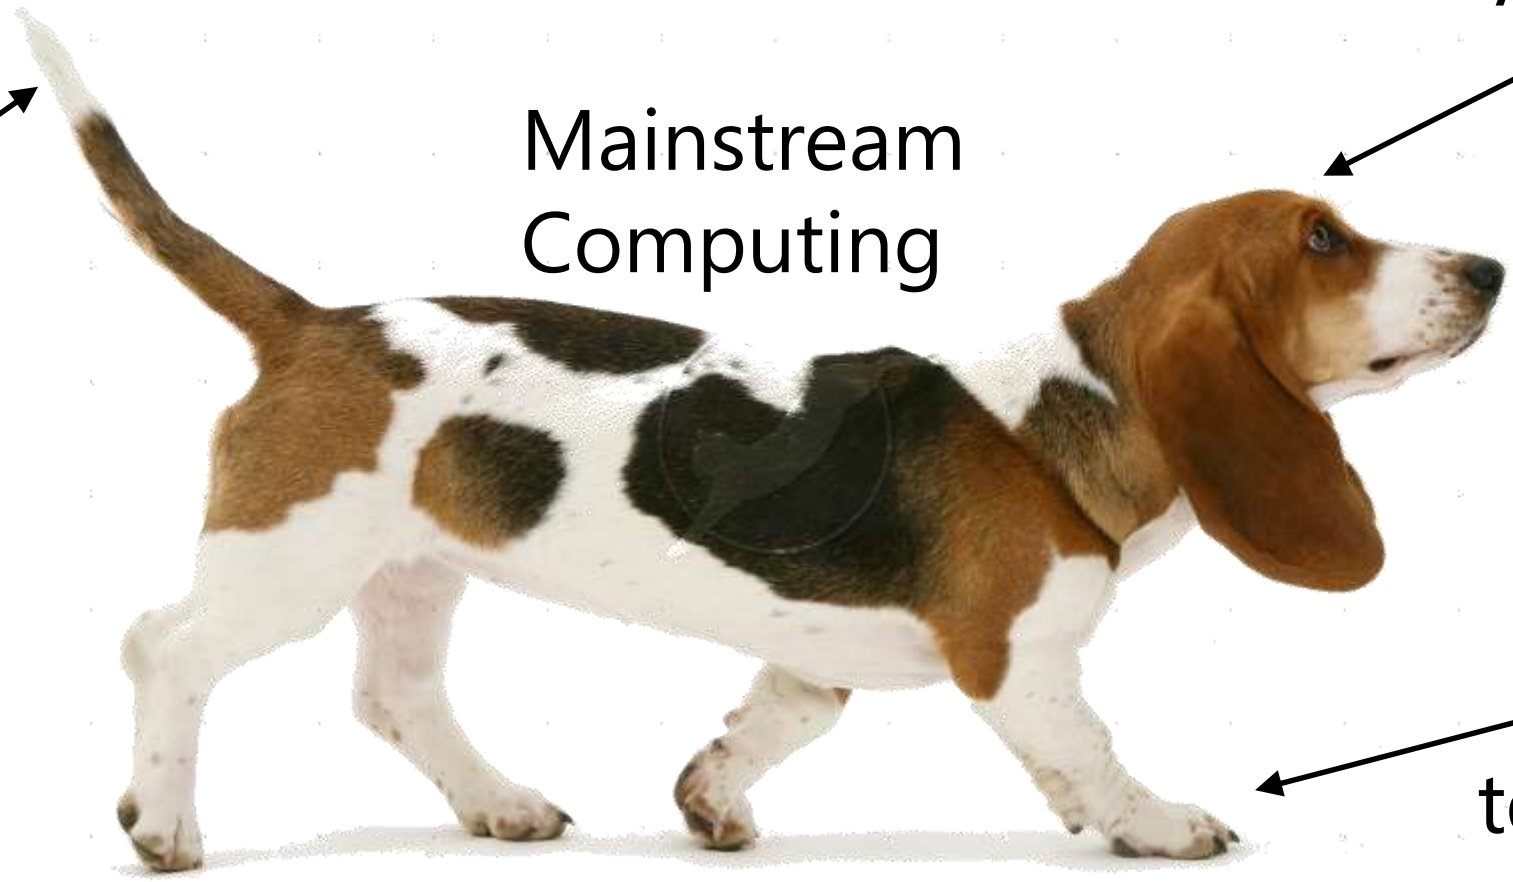

… or the money and the users may not follow you

Dot.com boom 2.0

Dominated by three areas

- Social networks
- Mobile apps
- Cloud services

Enabled by technology changes

- Higher speed broadband
- Smartphones and apps
- ***Machine learning and data analytics***

Clouds 101

It's a bird, it's a plane …

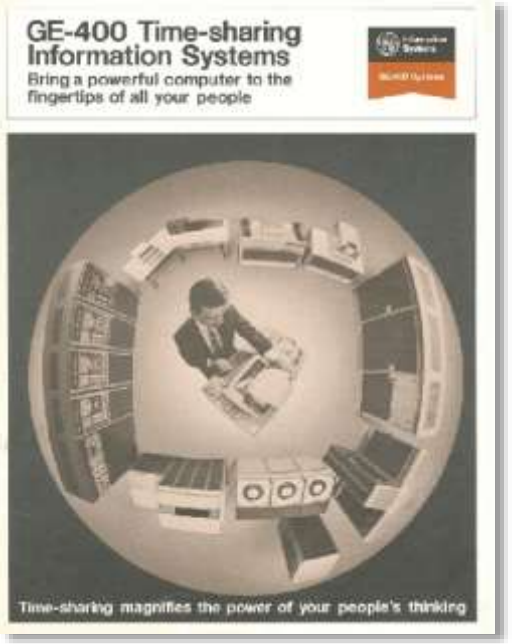

Some of us are old enough to remember public timesharing …

Data centers built for Internet services

- Search, email, e-commerce …

Microsoft Azure

amazon web services™

Google App Engine

Infrastructure optimization at scale

- Cloud software stacks, data centers, CDNs
- Virtualization, containers, SDNs, IaaS/PaaS/SaaS

mesosphere

kubernetes by Google

databricks

Rise of deep learning

- Big data and analytics

CoreOS

docker

Spark

THE UNIVERSITY OF IOWA

Cloud data centers: The iceberg view

It's all about ratios …

- Cost, performance, capacity

Jim Gray's four axes

- Networking
- Computation
- Storage
- Access

… and that is still driving change

- Modular building blocks
- Component integration
- Systemic resilience
- Cooling and power
- SDN and SDS, containers

Inspiration: Christian Belady

Google translate

skype™

Microsoft HoloLens

Natural Interfaces and Intelligent Assistants

Optimizing the continuum

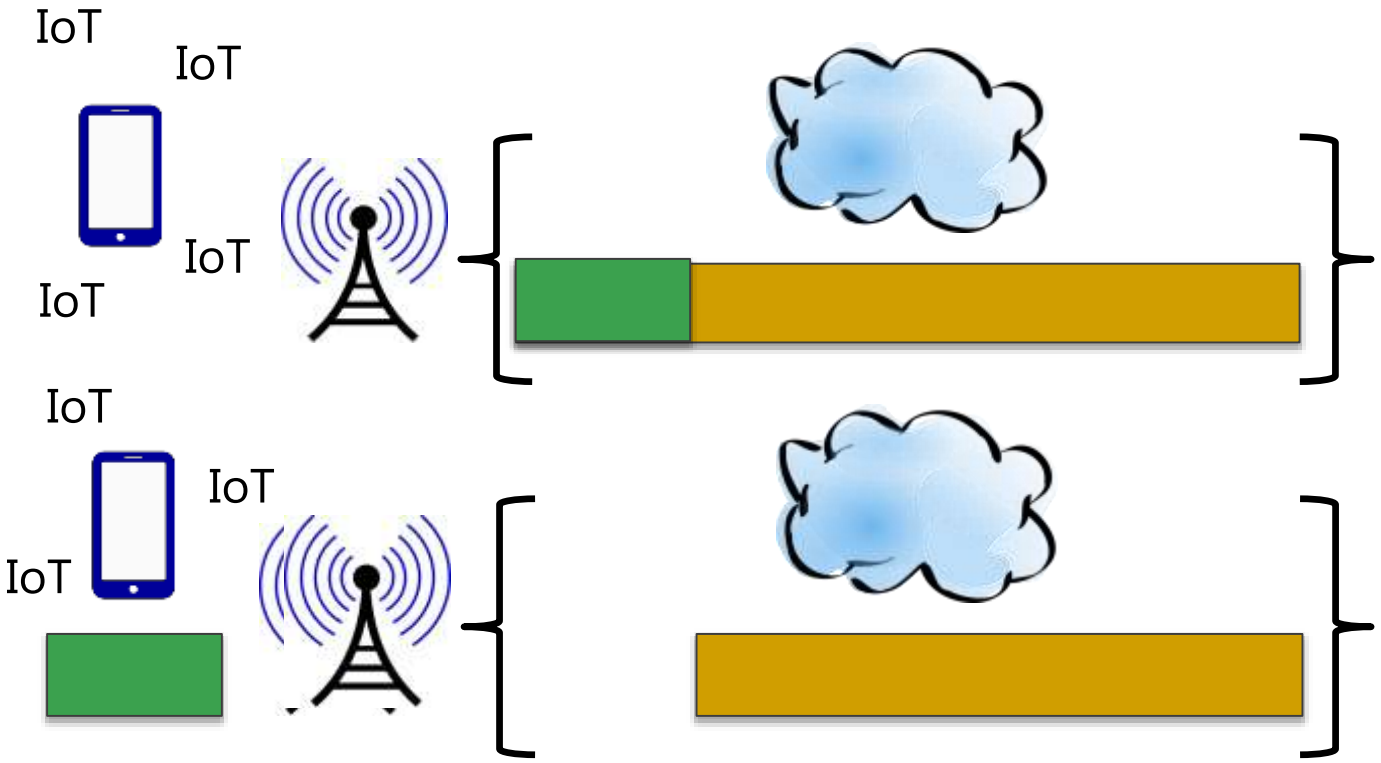

Latency
Bandwidth
Energy

Storage
Knowledge
Context

Think about augmented reality

Turtles (almost) all the way down

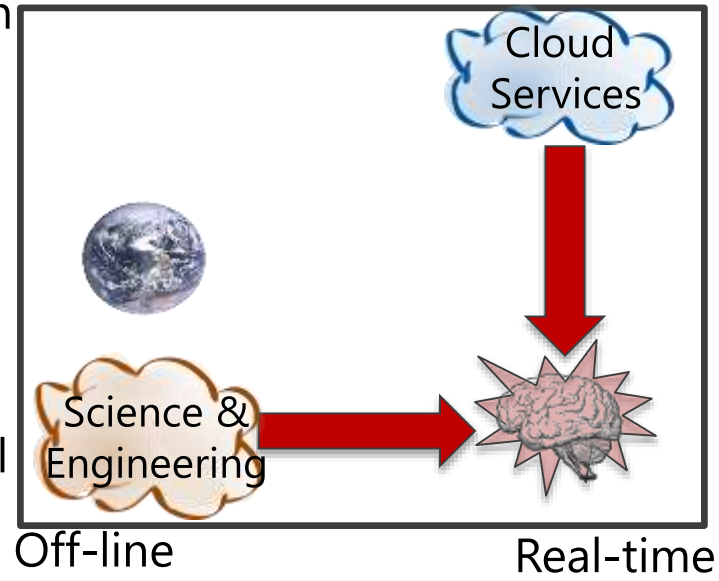

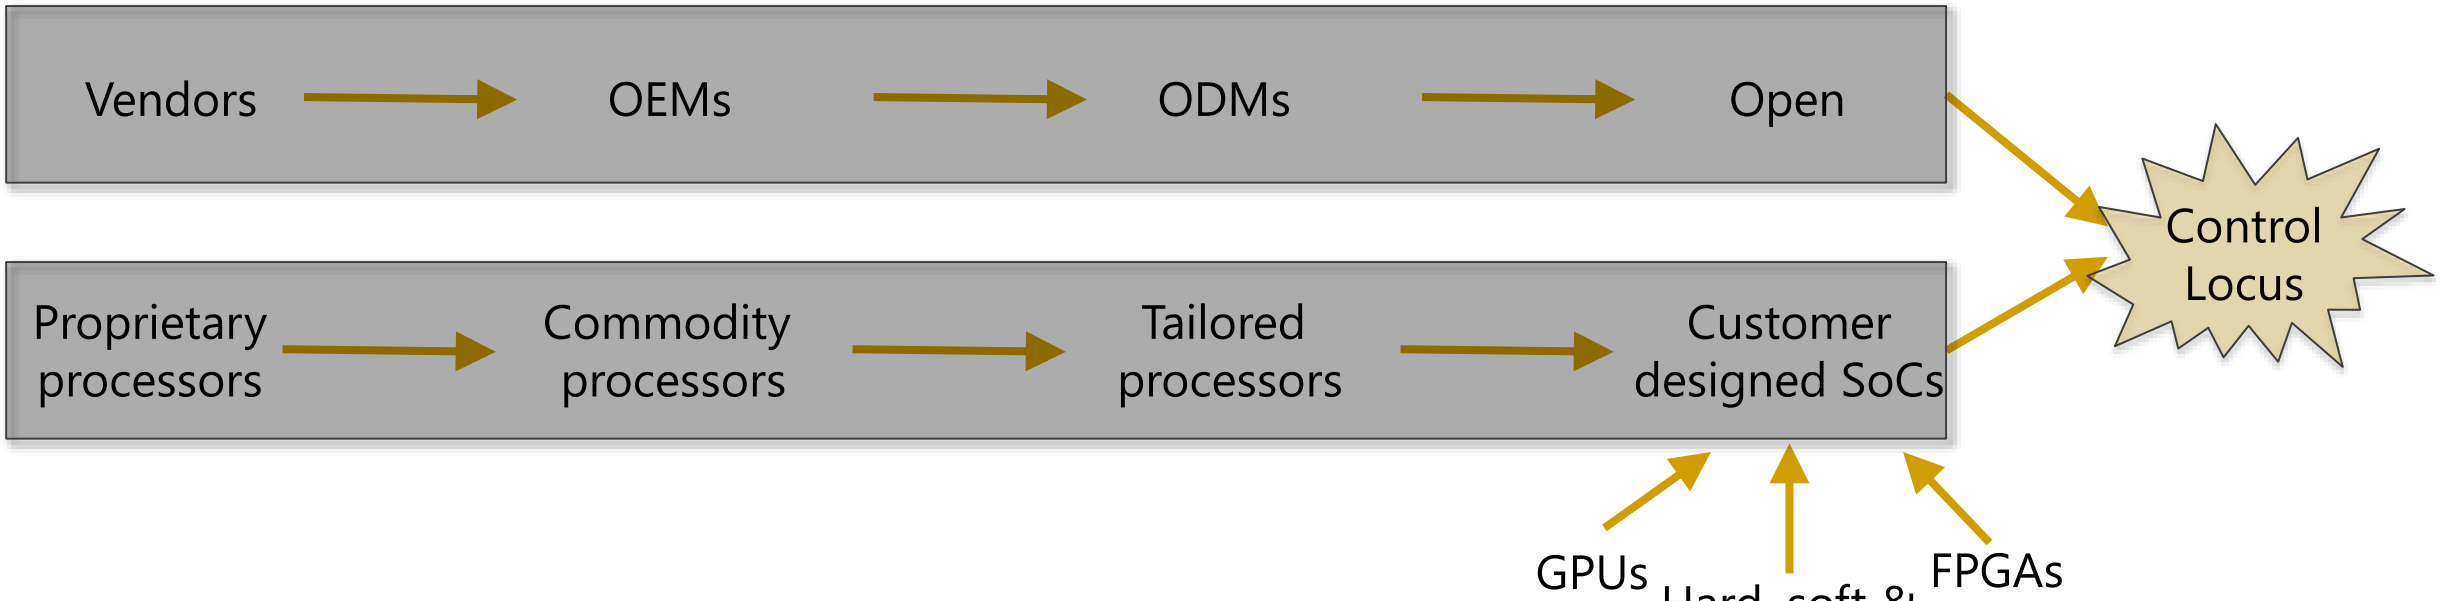

Software is eating the world – Marc Andreessen

Software is eating the world

Containerization (OS virtualization)

Software defined networks (SDNs)

- Merchant silicon revolution
- Network control
  - Uniform abstractions and separation of concerns
  - Control functions in OS not network switches

Software defined storage

- Abstraction and separation
- Storage virtualization
- Policy management

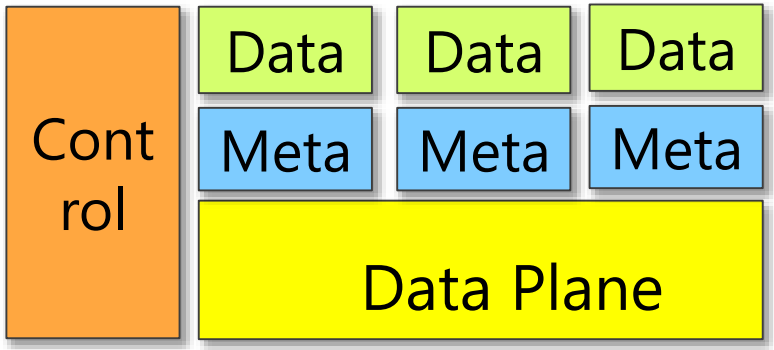

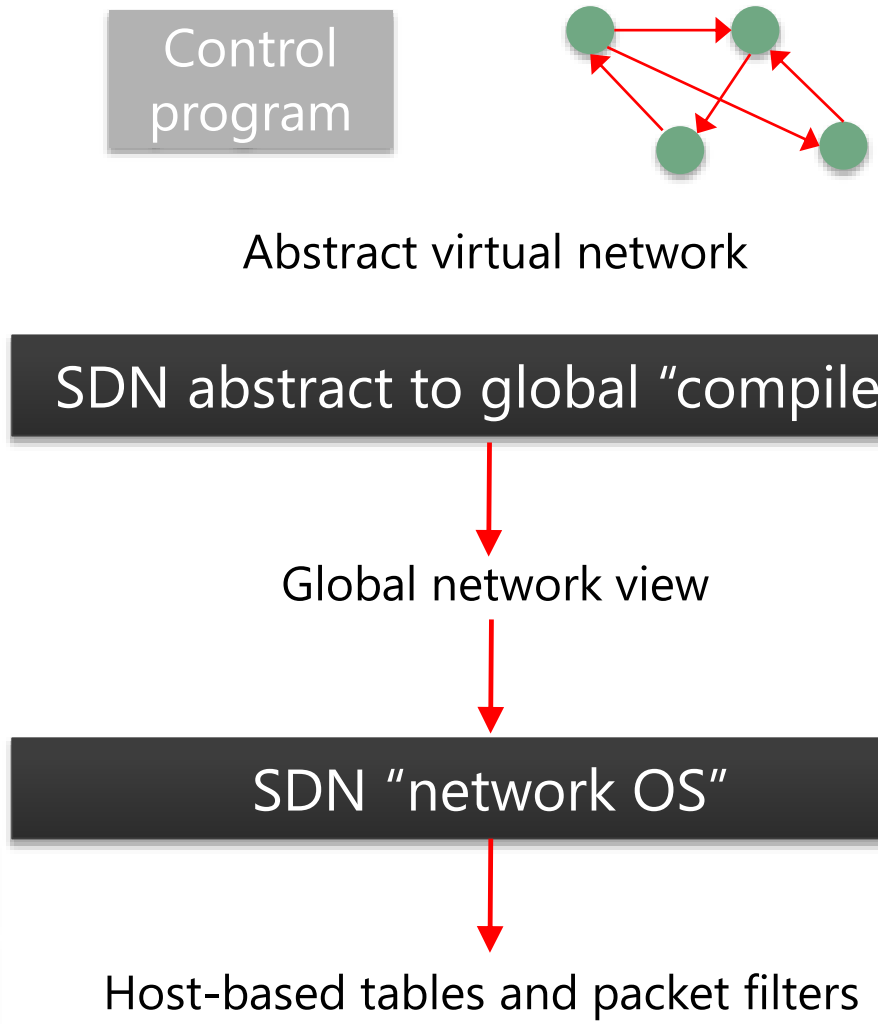

Acknowledgment: Dennis Gannon

Rebooting the IT Revolution

Energy efficient sensing and computing
Cyber-physical systems
Intelligent storage
Real-time communication ecosystem
Multilevel and scalable security
Next-generation manufacturing paradigm
Insight computing

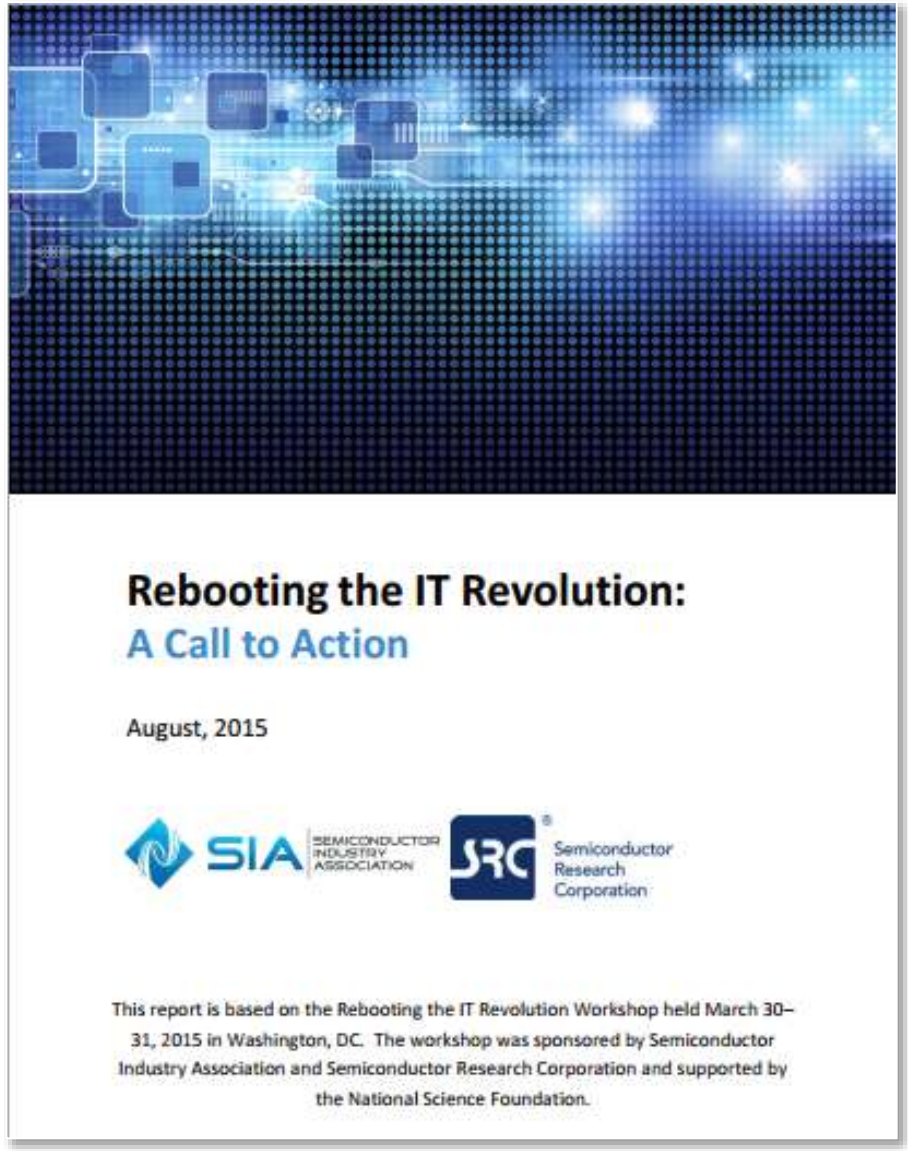

Computing energy efficiency

100 picojoules can

- Run a Cortex M0 for 10 cycles
- Write 1 bit of FLASH
- Write 300 bits of SRAM or DRAM
- Send 5 bits across LPDDR4
- Transmit 2 bits of UWB data
- Transmit 0.02 bits over Bluetooth LE

Remember

- 1 KW is 3.6 *megajoules*

We have to be more efficient

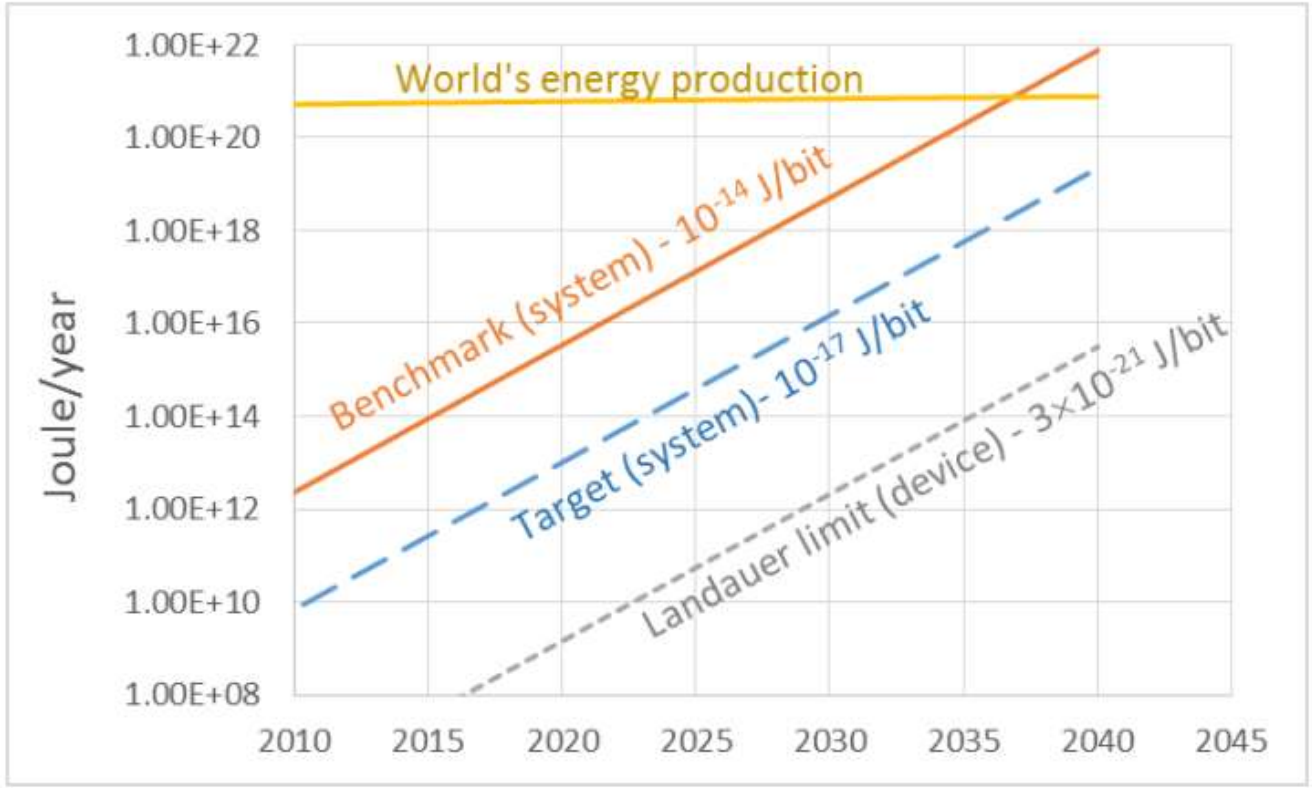

Sources: ARM, SIA/SRC (Rebooting the IT Revolution)

Analog brains and digital computers

Parallel distributed architecture	Mostly serial architecture
Low power (25W)	High power (100 MW)
Small volume (1 liter)	Large footprint (40M liters)
Asynchronous	Synchronous
Analog computing	Digital computing
Integrated memory/computation	Separated memory/computation
Intelligence via learning	Intelligence via programmed rules
Noisy components	Precise components
Low speed operation	High speed operation
Spontaneously active	Passive unless instructed

Source: SIA/SRC (Rebooting the IT Revolution)

Shifting computing perspectives
An error prone experiment

Accept inaccuracy/errors as the norm
- Verification and validation (V&V)

Computation as an {un}biased sample
- "An" approximate answer rather than "the" answer

Many axes
- Variable precision arithmetic
- Bit error tolerance
- Software correctness
- Energy adaptivity

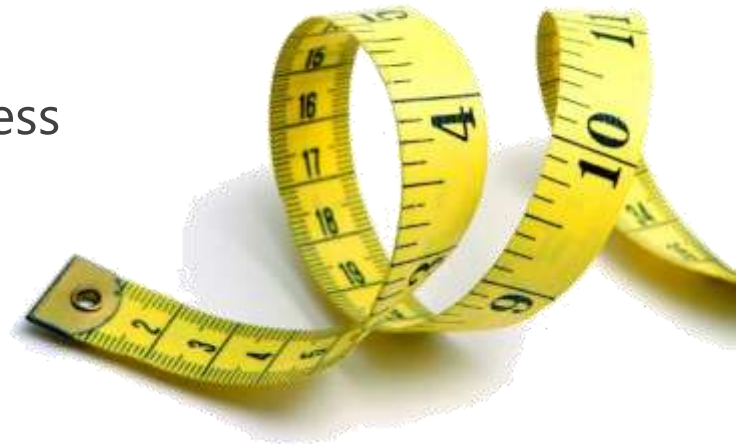

Analog, noisy inputs
- Sensors

Analog, error tolerant outputs
- Multimedia, data classification

Multiple acceptable answers
- Web search, climate models

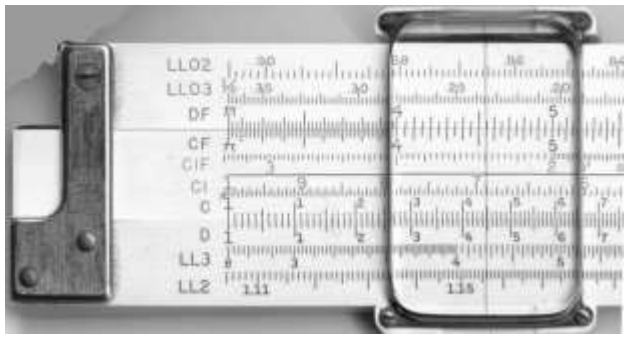

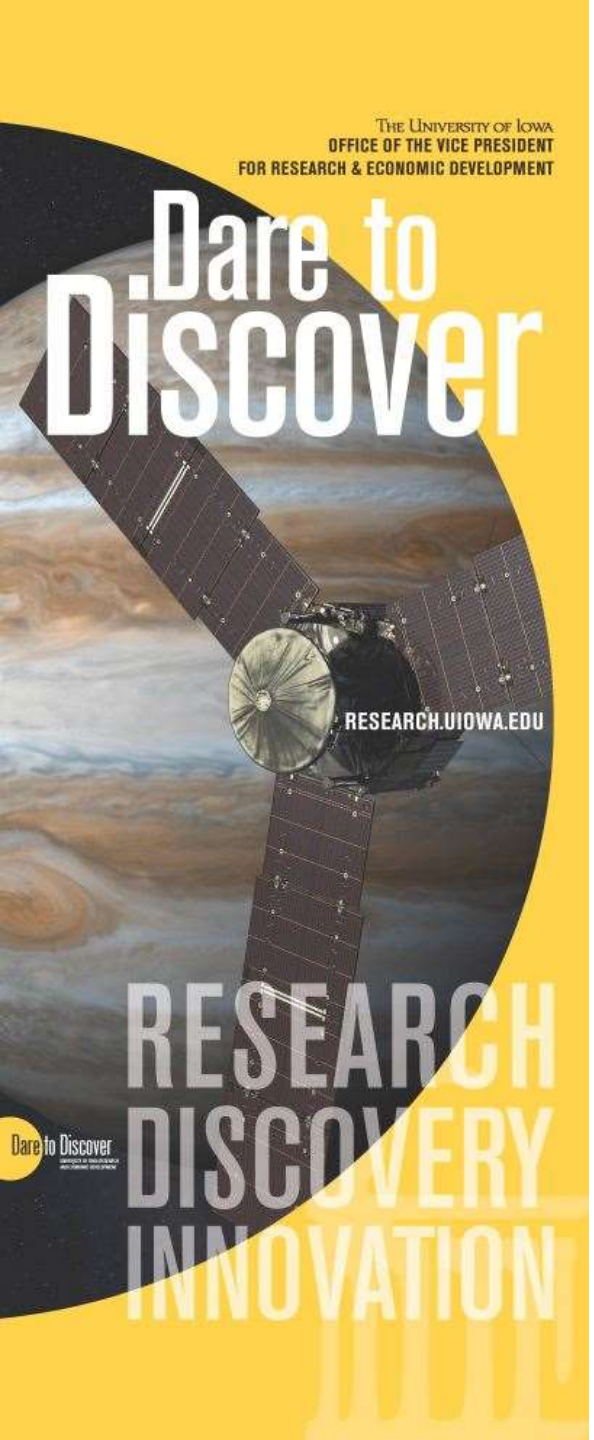

Thank You

The University of Iowa