Architectural Directions for Future Nanoscale Computing Systems

Kerry Bernstein
IBM T.J. Watson Research Center
Yorktown Heights, NY

MPSoC 2011
4 July, 2011
Beaune, France

1927 Conference at Solvay

The Contributions of 4 Physicists in high speed computing

- Max Planck - Energy come in discrete packets ($h\mu$)
- Niels Bohr - H_2 Model
- Erwin Schroedinger - Expansion of the Shell Model / Wave Eq
- Wolfgang Pauli - Modern Energy Bands and Exclusion

Who Let the Dogs Out?

Engineers' scaling of CMOS is confronting fundamental limitations.
To extend CMOS, industry must "virtually scale" with structure, materials

The Engine Driving The Digital Age

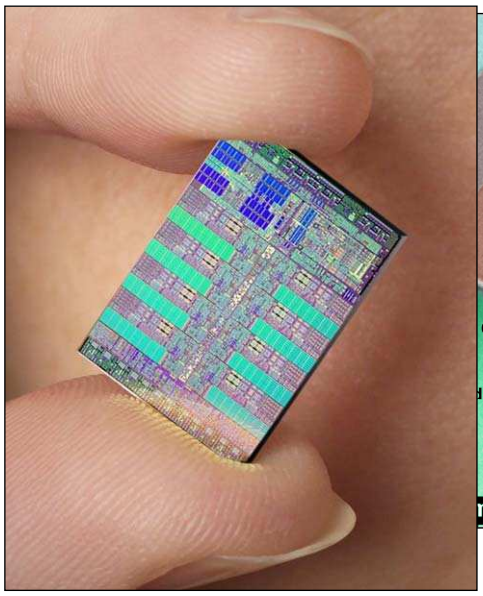

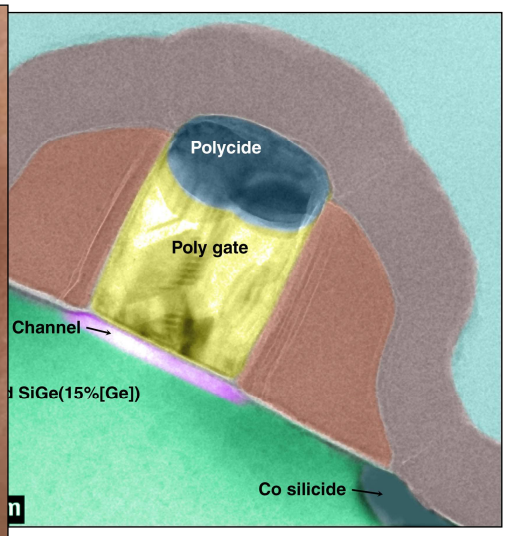

IBM's "Cell" 8-core Multiprocessor and
the underlying fundamental transistor

Level-set: Workloads

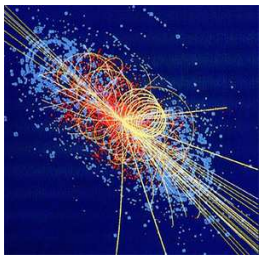

- **Scientific**
  - Highly regular, predictable patterns allow streaming data, cache to λP
  - Performance is directly proportional to BW
  - High Bus Utilization
  - Low Miss Rate

- **Commercial**
  - Unpredictable irregular patterns
  - Requires low bus utilization to avoid clogs in the event of a burst of misses (< 30% bus utilization)
  - Miss rate follows Poisson process (random)
- **Client/Server Model or Workstation/PC Model?**
- **Both spaces need BW, but use it very differently**

Level-set: Relative FLOP values

1 Megaflop	**2004**	**"Superbabies: Baby Geniuses 2"**
2.4 Gigaflops	**1999**	**AMD Athlon at 600 MHz**
11.38 Gigaflops	**1997**	**Deep Blue, IBM Chess Engine**
115 Gigaflops	**2005**	**XBOX360 with XENOS Graphics**
80 Teraflops	**2010**	**Jeopardy, 5 IBM Power-7 racks**
2.507 Petaflops	**2010**	**Fastest Supercomputer, Natl Univ. of Defense Tech.**
100 Petaflops	**200,000 BC**	**Human Brain (est)**

Computer Performance	
Name	**FLOPS**
yottaFLOPS	10^{24}
zettaFLOPS	10^{21}
exaFLOPS	10^{18}
petaFLOPS	10^{15}
teraFLOPS	10^{12}
gigaFLOPS	10^{9}
megaFLOPS	10^{6}
kiloFLOPS	10^{3}

"Why Explore New Compute Architectures?"

a) The desire to solve a range of problems for which no clear algorithmic solution exists

b) The need for practical solutions to problems for which silicon offers only the most cumbersome of solutions

c) The impending predicted limitations in existing silicon technologies

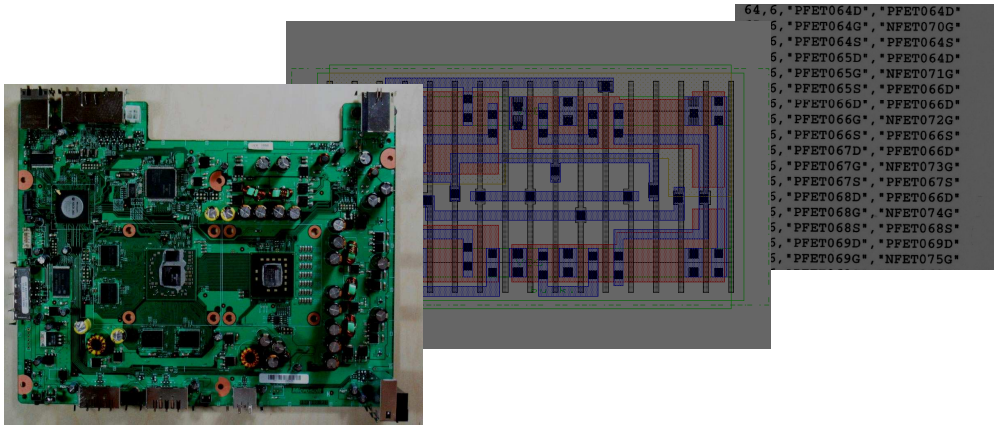

Limits to the Present Technology

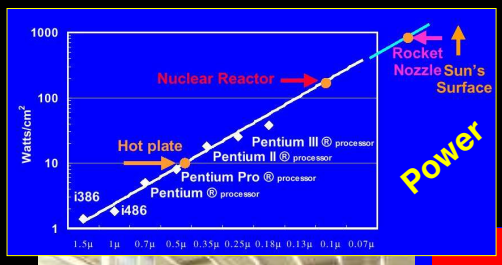

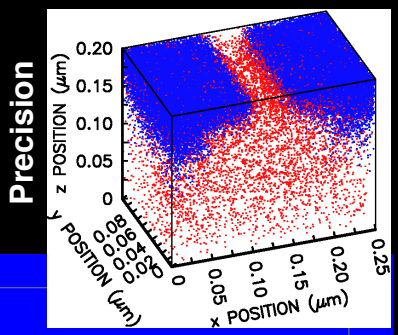

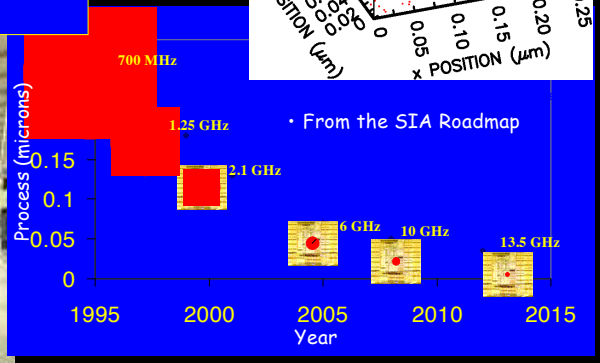

The Nanoelectronics Research Initiative (NRI)

sponsored by the Semiconductor Research Corporation (SRC)

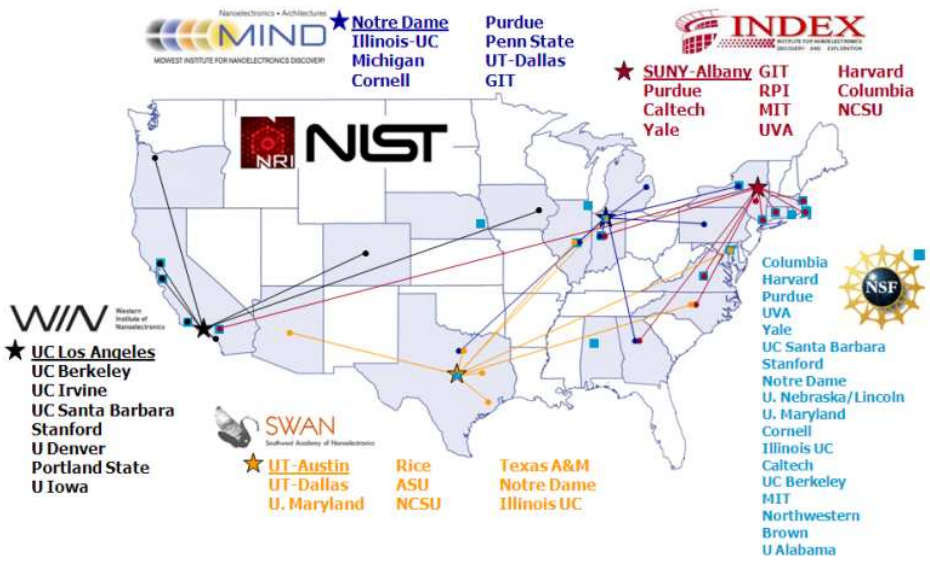

Thermodynamic vs Information Entropy

- VLSI is confronting atomistic limitations to scaling; soon these limitations will be quantum-mechanical
- Scaling CMOS requires the large investments of energy to assert structure where structure does not wish to exist
- VLSI as well as the Universe wishes to devolve to a state of lowest energy, highest disorder.
- The post-CMOS mission:
  Find those materials, structures, architectures which extend transaction retirement rates, and simultaneously improve cost, energy use, and reliability.

Raising the Bar for New Switches

An Evolution of Demands on the Next Device

1983 (Keyes)
- Inversion
- Logic gain
- Self Restoring
- High Ion/ Ioff

2001 (Svensson)
- Flexibility
- Isolation
- Low cost
- Reliability
- Speed

2010 (NRI Architecture Study)
- High Complexity Content
- Low Logical Effort
- Low Energy-Delay Product
- Security "Root-of-Trust"
- High Radiation Immunity

Essential Architectural Features for Any Switch Succeeding CMOS

Computational Efficiency
New switches must be organized in architectures which extend the computation completed per unit of energy, area, and device

CyberSecurity
New compute structures will provide hardware-based root-of-trust: i.e. tamper detection, watermarking, and side-channel attack resilience

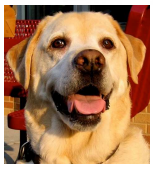

Reliability
Fully servicable post-CMOS devices will assure serviceability and reliability with self healing, self-reconfigurable capabilities

Complexity Accommodation
We'll need the ability to describe increasingly complex workloads to processors with exponentially increasing machine state count

Exposure	Present Mitigation
Local/Network Malicious Code	
Trojan	Antivirus Software
Virus / Worm	Antivirus Software
Local/Network Malicious Data	
Data directed attacks	IDS/IPS
Buggy code/firmware	Static/Dynamic Analysis
Physical Attack	
Supply Chain	Sampling
Local	Armed guards
Platform	
Covert/Side channel Attack	Sampling
Emissions	Faraday cage
Design Flaws	Formal Methods
Broken Cipher	Design Metholodogy
Counterfeit Hardware, Trojans	Monitoring Watermarking
Operation	
Untrusted Operator – Intentional	Screening
Trusted Operator – Inadvertent	Human Factors
Process / Operator Failure	Log analysis tools
Configuration Failure	Monitoring tools

"The Knowledge" and Compute Plasticity

Self-organization, Self reconfigurability a hallmark of intelligence?

Unreliable Components

- SWCNTs and Chirality

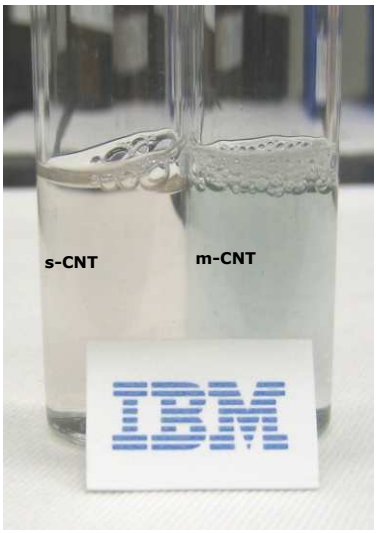

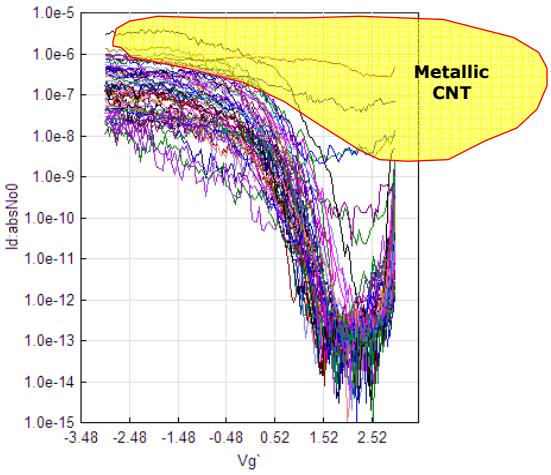

A. The Notion of Logical Effort (LE)

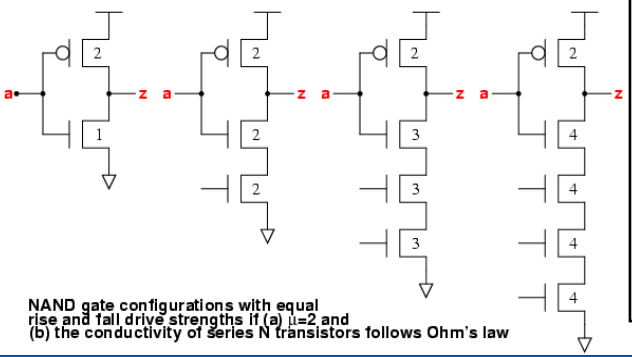

NAND gate configurations with equal rise and fall drive strengths if (a) μ=2 and (b) the conductivity of series N transistors follows Ohm's law

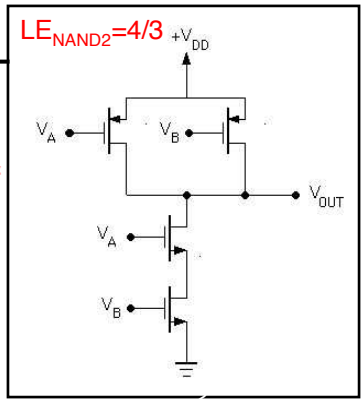

LE describes effect of (ckt topo + dev tech) upon the ability to produce output

Within CMOS, Logical Effort is simply:

"The number of times worse a given circuit is at driving an output load, compared to a simple inverter with the same amount of input capacitance."

Alternately:

"The "ratio of a given circuit input's capacitance to that of an inverter delivering the same output current."

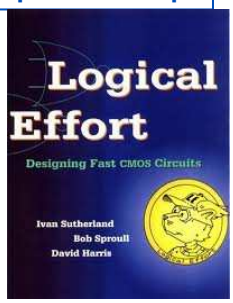

Processor Complexity Space*

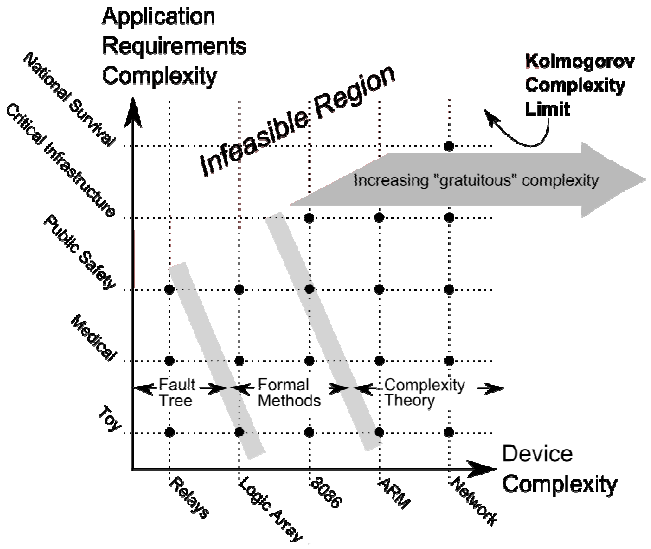

Increased complexity will necessary to accommodate new function we ask of uPs

Chip Complexity offloads for economy, ie. software, compiler work is shifted instead to hardware

Designer Population or EDA must scale with Moore's Law unless more complexity is captured in the switch

With complexity comes a greater ability to hide machine states. Only about 1/3 used are presently used.

*After Jackson Mayo
Sandia Natl Labs

Unconventional Structures Realizing Conventional Operations

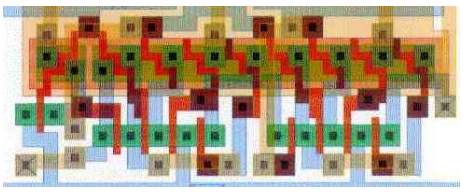

CMOS Inverter

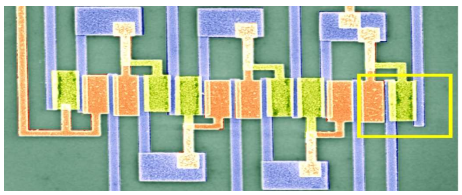

CNT Inverter
(J. Appenzeller
Purdue University)

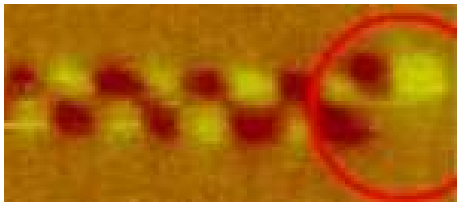

NML Inverter
(M. Niemier,
Univ. of Notre Dame)

Exemplary Structures and information tokens executing logical inversion

Raw Data

NAND2 Delay-Energy-Area Space

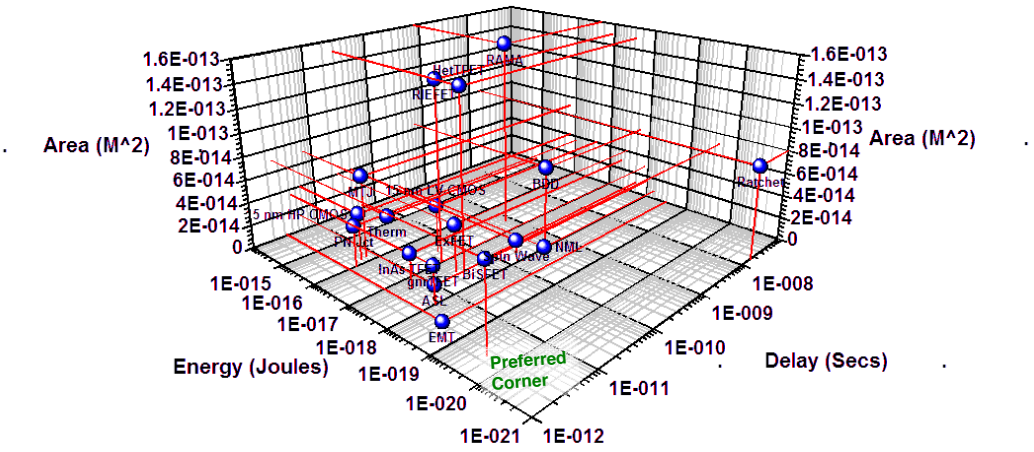

- **Convergence noted in data reported.**
- **Data cluster of Many-Spin and Few Spin Structures**
- **5 energy decades for approx equivalent combinatorial delay**
- **Very little area differentiation**

NAND2 Delay vs Energy

Raw Data

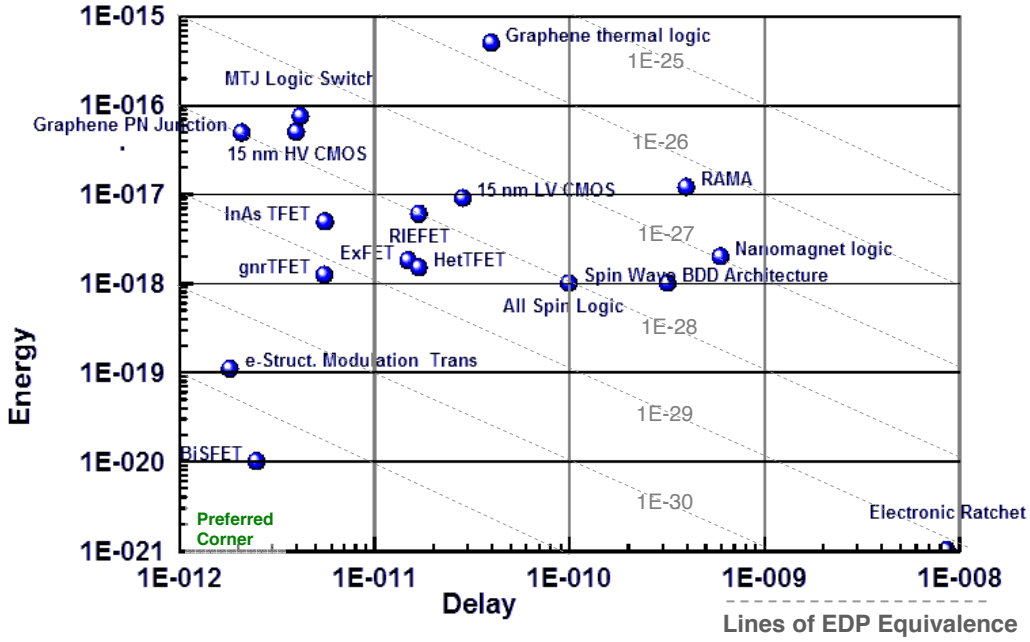

A potential Delay-Energy minima exists at approximately 1E-29

Ultimate Boundaries: Minimum Switching Energy

From "The Physical Limits of Computing" M.P. Frank, *Computing in Science and Eng*, Oct, 2001

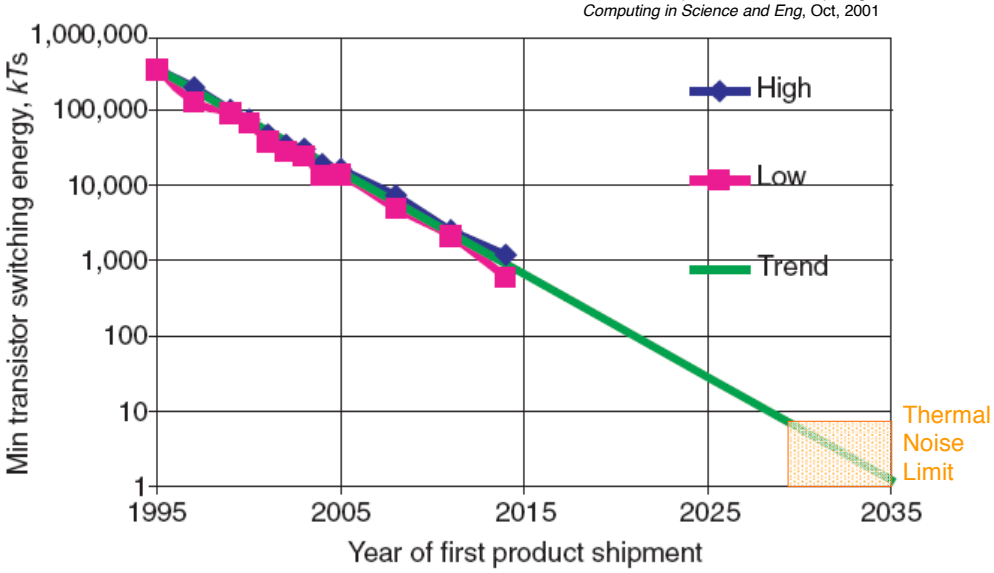

Logic at the Unit Gate Level

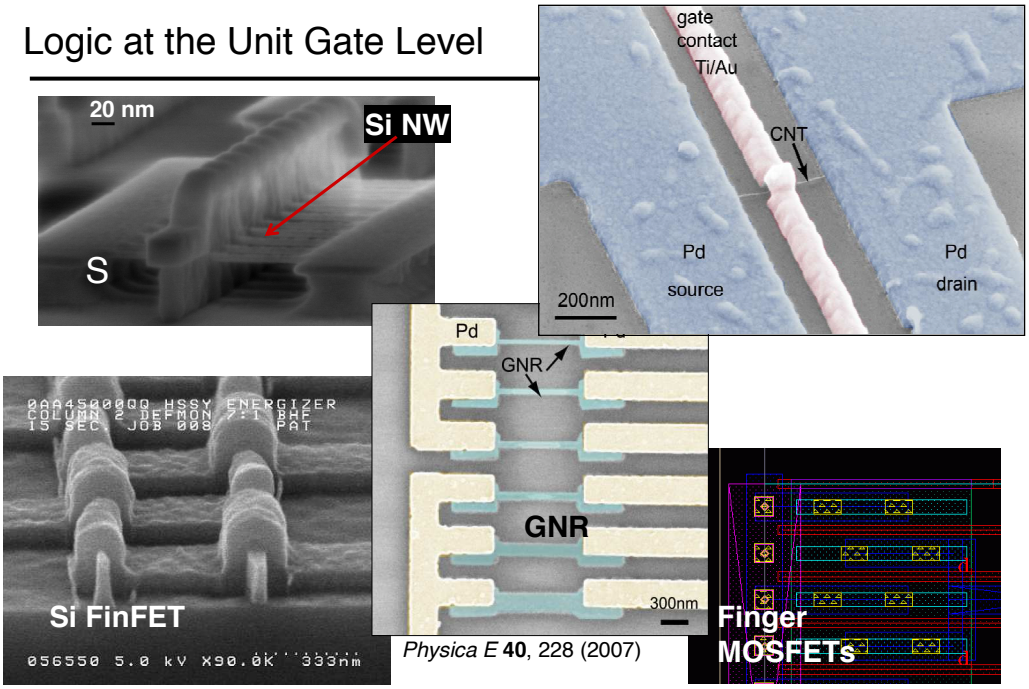

Physica E **40**, 228 (2007)

Switches going forward *naturally quantize their output* – can we exploit this in new architectures?

Interconnect "Span of Control

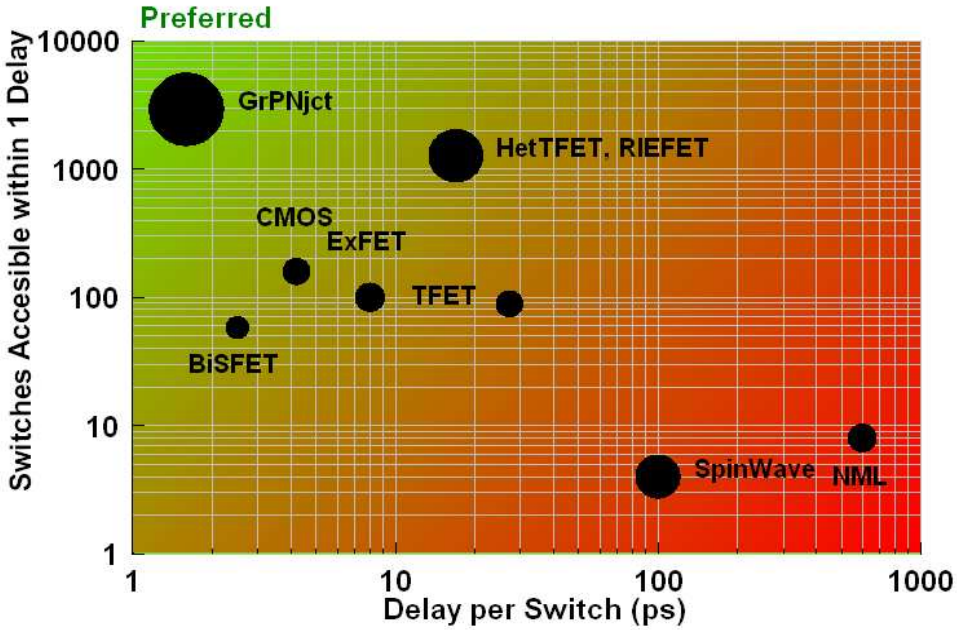

Evolution of 3D Integration

Technology Investment in the Z-Dimension

- 3D Technologies continue the sequence of interconnect advances
- Return balance to device scaling
- Enable new capabilities not available in 2D

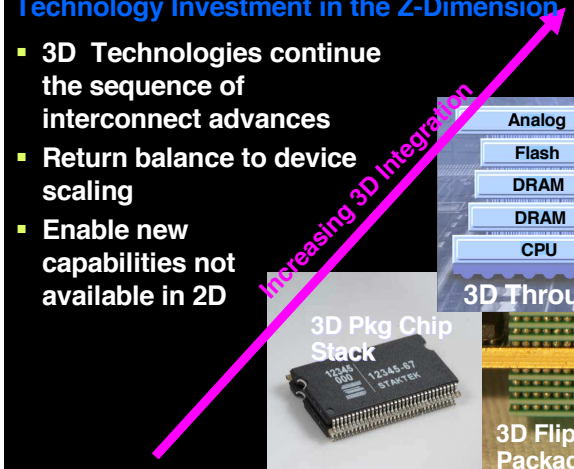

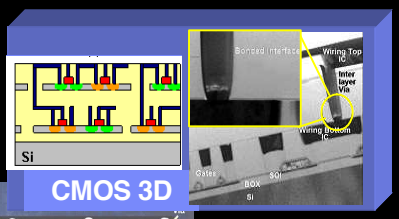

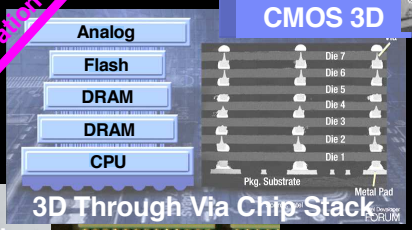

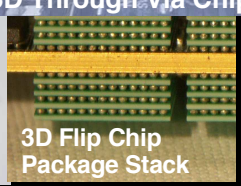

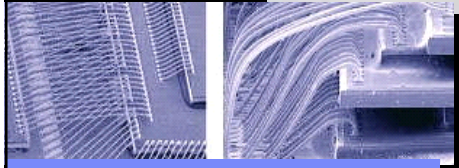

- 3D Packaging R&D now pervasive in industry, academia
- Through-via technology emerging as predominant path
- 3D has *always* been large volume, but now integrating higher technologies

Trends in Architecture Performance

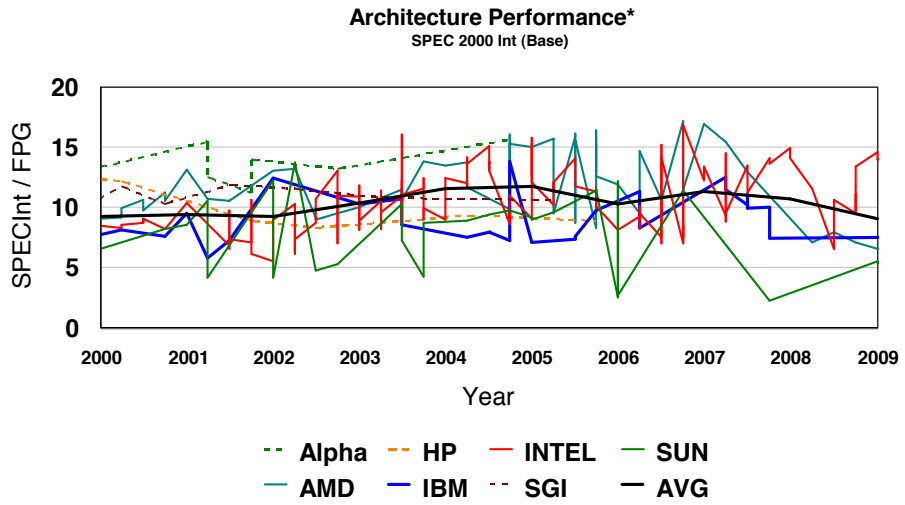

Highest reported SPEC2000 INT per (adj)FPG Generation

FPG, SPECmark approximated when necessary; Broken line = discontinued series

K. Bernstein 11/06

Exemplary Alternative Architectures

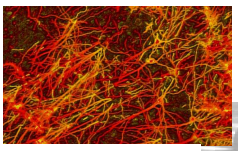

Neuromorphic

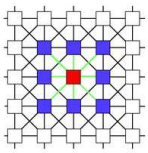

Cellular
Non-Linear Nets

Superposition and
Entanglement

John
Von Neumann

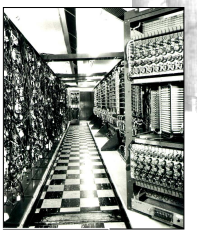

Analog

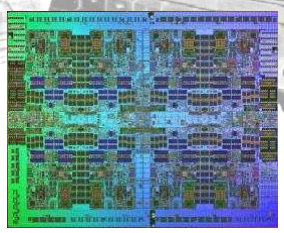

Many-core TLP λProc

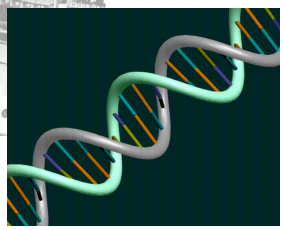

Algorithmic
Self-Assembly

Garry Kasparov vs IBM Deep Blue (1997)

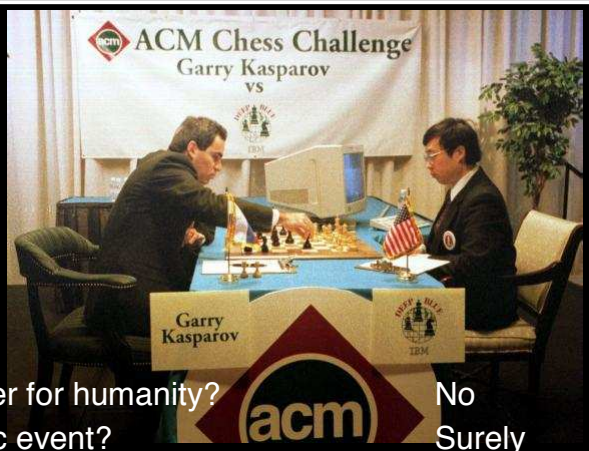

- A disaster for humanity? No
- A historic event? Surely
- Can Deep Blue think? NO !! (Not yet)
- Has our perspective changed? Perhaps ….

IBM "WATSON" (2011)
Common Language Question Answering (QA) System

Components of Processor Performance

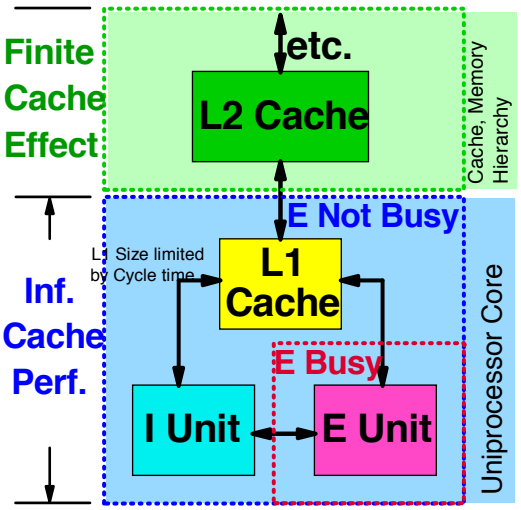

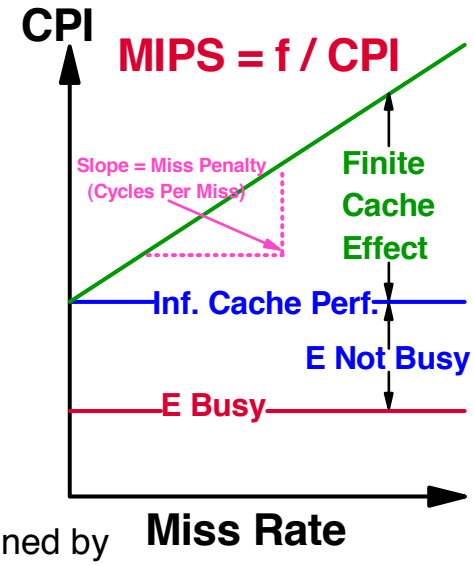

Delay is sequentially determined by
a) ideal processor,
b) access to local cache, and
c) refill of cache

Virtualization

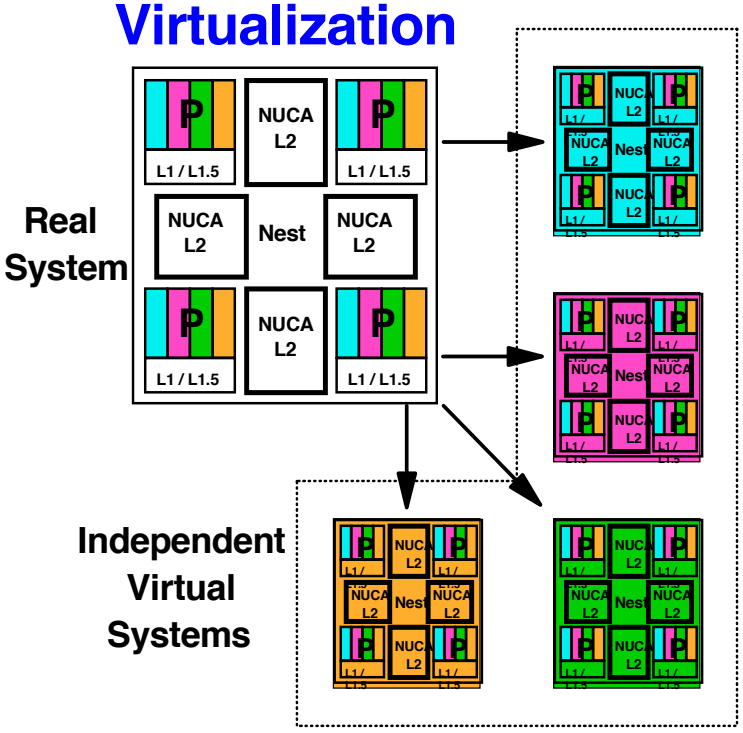

"Looks like"
4 independent systems, each with 16 cores!

Intermiss Distance Density

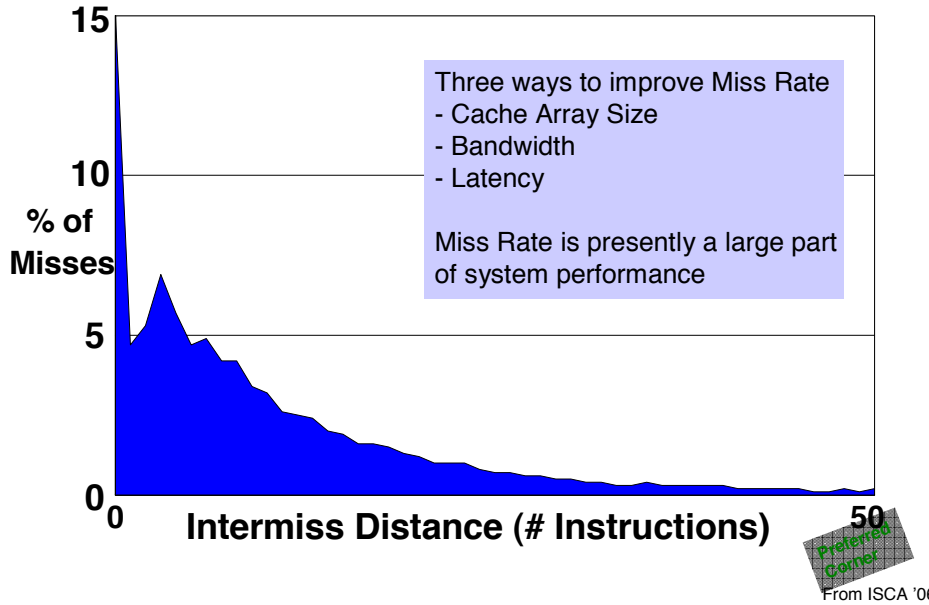

From ISCA '06 Keynote address by Phil Emma, IBM

The Rising Waters of Bio-emulation

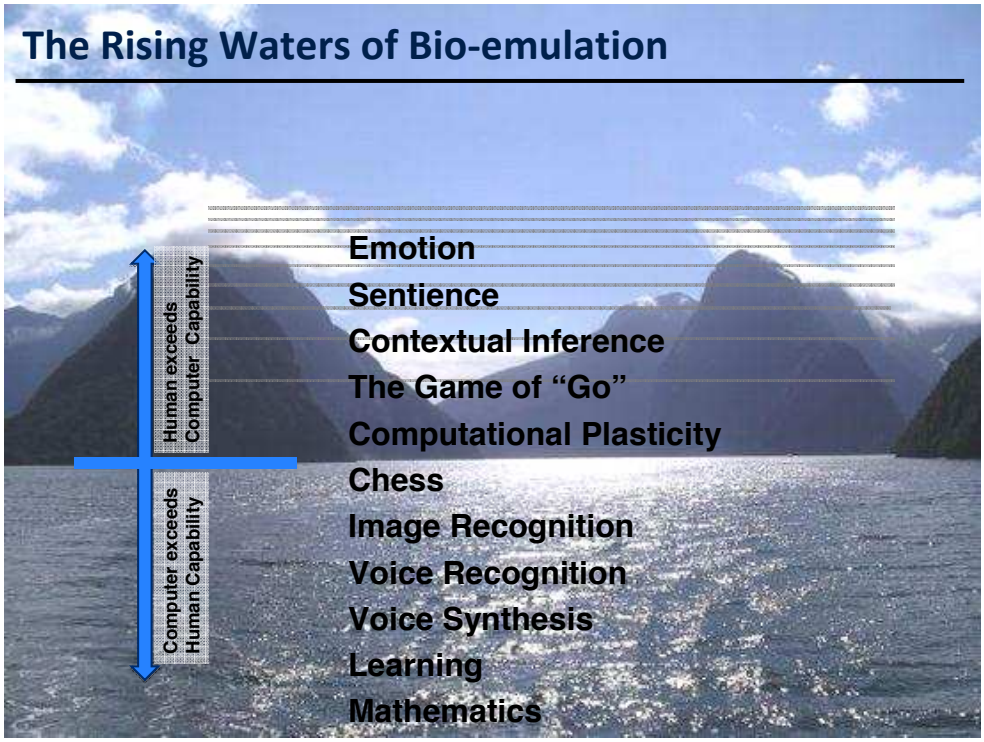

Scaling & Machine "Intelligence"

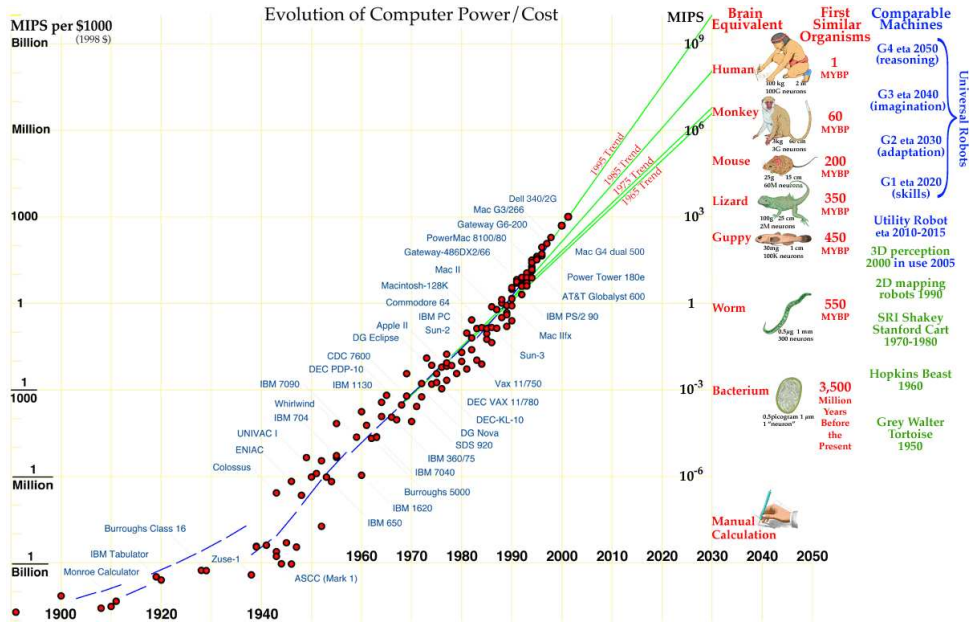

From *Robot: Mere Machine to Transcendent Mind*, Hans Moravec, Oxford University Press, 1999

Summary

- HPC is confronting profound limitations to continued scaling.
- Next-gen logic switches haven't yet surpassed CMOS in energy, delay, and area simultaneously; in other architectures these new devices may already be superior.
- Transport mechanisms will have a profound challenge communicating new information tokens effectively.
- [(low-energy) = (high-delay)] conundrum continues.
- New architectures will need to sustain higher complexity, security, reliability, efficiency.
- Post-CMOS architectures succeed with our own species' ability to think in parallel.
- Whole new realms of compute capability await the introduction of the next switch and its organization.

Medieval Nanoscience

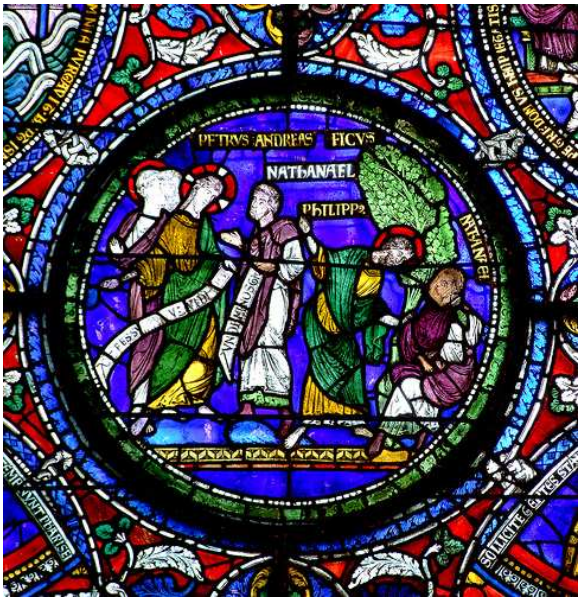

Metallic nano-particles exhibit radiant colors due to plasmon resonance, Metals/Oxides cause the strong pure colors of medieval stained glass windows; gold creates deep red; copper, blue; iron, green. Plasmons on the particle surface move to absorb blue and yellow light but reflect longer κ red, providing the characteristic ruby color

Late 12th century.

North window, northern aisle, choir section. Canterbury Cathedral, Kent, England