SAND2005-1079C

Sandia Zettaflops Story

A Million Petaflops

Erik P. DeBenedictis

Sandia National Laboratories

February 24, 2005

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Applications and $100M Supercomputers

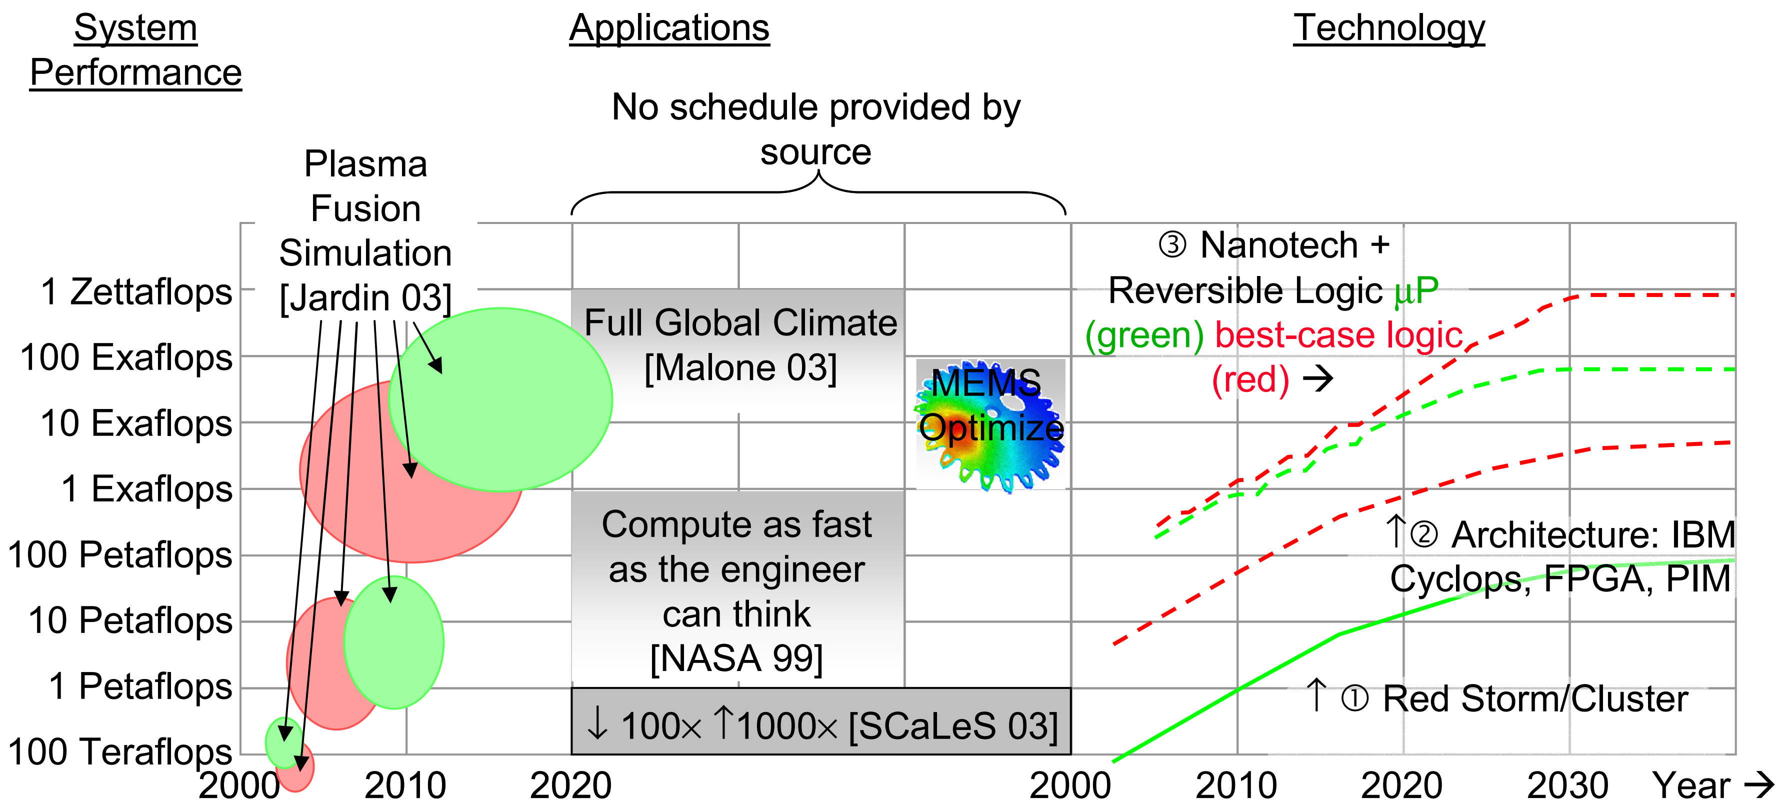

[Jardin 03] S.C. Jardin, "Plasma Science Contribution to the SCaLeS Report," Princeton Plasma Physics Laboratory, PPPL-3879 UC-70, available on Internet.
[Malone 03] Robert C. Malone, John B. Drake, Philip W. Jones, Douglas A. Rotman, "High-End Computing in Climate Modeling," contribution to SCaLeS report.
[NASA 99] R. T. Biedron, P. Mehrotra, M. L. Nelson, F. S. Preston, J. J. Rehder, J. L. Rogers, D. H. Rudy, J. Sobieski, and O. O. Storaasli, "Compute as Fast as the Engineers Can Think!" NASA/TM-1999-209715, available on Internet.
[SCaLeS 03] Workshop on the Science Case for Large-scale Simulation, June 24-25, proceedings on Internet a http://www.pnl.gov/scales/.
[DeBenedictis 04], Erik P. DeBenedictis, "Matching Supercomputing to Progress in Science," July 2004. Presentation at Lawrence Berkeley National Laboratory, also published as Sandia National Laboratories SAND report SAND2004-3333P. Sandia technical reports are available by going to http://www.sandia.gov and accessing the technical library.

Sandia National Laboratories

Outline

- An Exemplary Zettaflops Problem
- The Limits of Moore's Law
- Beyond Moore's Law
  - Industry's Plans
  - Nanotech and Reversible Logic
- Conclusions

FLOPS Increases for Global Climate

- 1 Zettaflops
- 100 Exaflops
- 1 Exaflops
- 10 Petaflops
- 100 Teraflops
- 10 Gigaflops

Issue	Scaling
Ensembles, scenarios $10\times$	Embarrassingly Parallel
Run length $100\times$	Longer Running Time
New parameterizations $100\times$	More Complex Physics
Model Completeness $100\times$	More Complex Physics
Spatial Resolution $10^4\times$ ($10^3\times$-$10^5\times$)	Resolution
Clusters Now In Use (100 nodes, 5% efficient)	

Ref. "High-End Computing in Climate Modeling," Robert C. Malone, LANL, John B. Drake, ORNL, Philip W. Jones, LANL, and Douglas A. Rotman, LLNL (2004)

Sandia National Laboratories

Exemplary Exa- and Zetta-Scale Simulations

- Sandia MESA facility using MEMS for weapons
- Heat flow in MEMS not diffusion; use DSMC for phonons
- Shutter needs 10 → Exaflops on an overnight run for steady state
- Geometry optimization → 100 Exaflops overnight run
  - Adjust spoke width for high b/w no melting

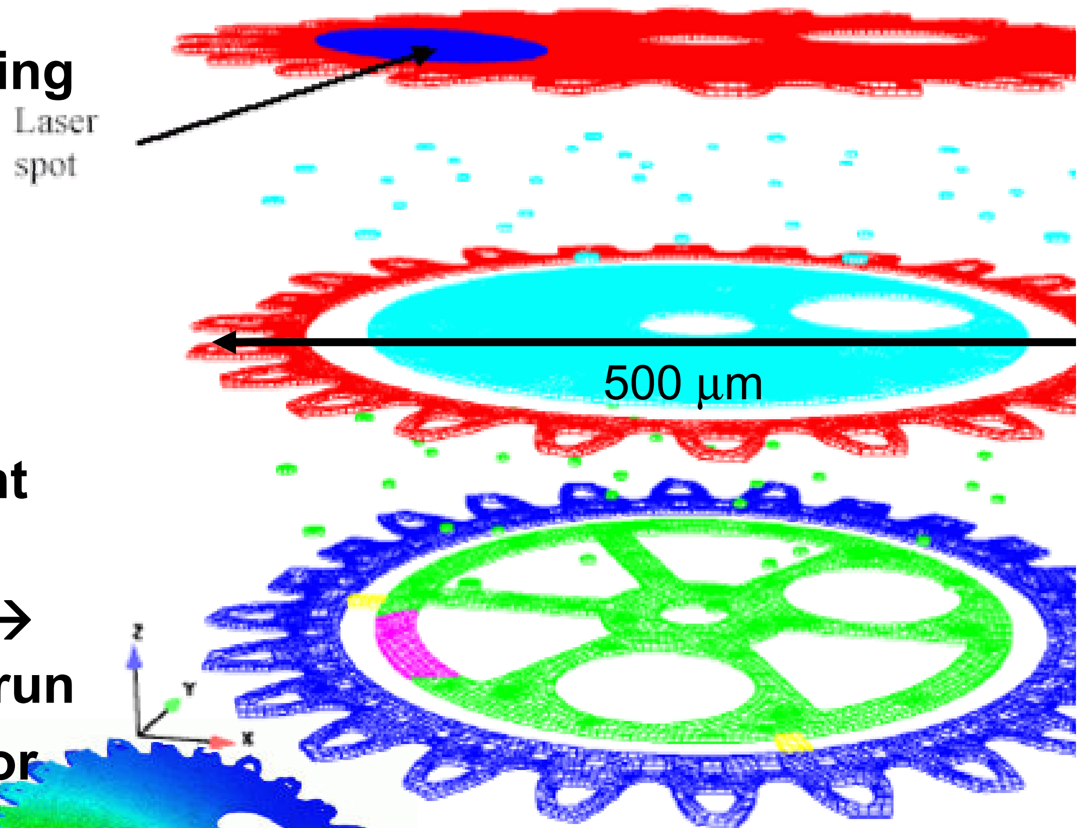

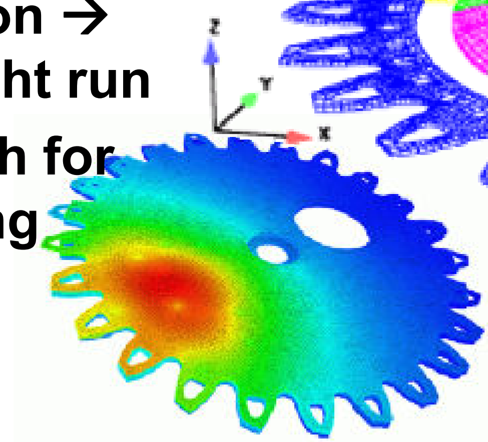

Sandia National Laboratories

FLOPS Increases for MEMS

100 Exaflops

10 Exaflops

30 Petaflops

600 Gigaflops

5 Gigaflops

Issue	Scaling
Optimize 10×	Sequential
Run length 300×	Longer Running Time
Scale to $500\mu m^2 \times 12\mu m$ disk 50,000×	Size
2D → 3D 120×	Size
2µm×.5µm×3µs 2D film 10 × 1.2 GHz PIII	

Sandia National Laboratories

Outline

- An Exemplary Zettaflops Problem
- The Limits of Moore's Law
- Beyond Moore's Law
  - Industry's Plans
  - Nanotech and Reversible Logic
- Conclusions

*** This is a Preview ***

Best-Case Logic	Microprocessor Architecture
2×10^{24} logic ops/s	
100 Exaflops	800 Petaflops
← 125:1 →	
25 Exaflops	200 Petaflops
4 Exaflops	32 Petaflops
1 Exaflops	8 Petaflops
	80 Teraflops
	40 Teraflops

Expert Opinion: 100 Exaflops – 800 Petaflops (← 125:1 →)

Estimate: 25 Exaflops – 200 Petaflops; 4 Exaflops – 32 Petaflops

Physical Factor	Source of Authority
Reliability limit 750KW/($80k_BT$)	Esteemed physicists (T=60°C junction temperature)
Derate 20,000 convert logic ops to floating point	Floating point engineering (64 bit precision)
Derate for manufacturing margin (4×)	Estimate
Uncertainty (6×)	Gap in chart
Improved devices (4×)	Estimate
Projected ITRS improvement to 22 nm (100×)	ITRS committee of experts
Lower supply voltage (2×)	ITRS committee of experts
Red Storm	contract

Assumption: Supercomputer is size & cost of Red Storm: US$100M budget; consumes 2 MW wall power; 750 KW to active components

Metaphor: FM Radio on Trip to Austin

- You drive to Austin listening to FM radio
- Music clear for a while, but noise creeps in and then overtakes music

- Analogy: You live out the next dozen years buying PCs every couple years
- PCs keep getting faster
  – clock rate increases
  – fan gets bigger
  – won't go on forever
- Why…see next slide

Details: Erik DeBenedictis, "Taking ASCI Supercomputing to the End Game," SAND2004-0959

FM Radio and End of Moore's Law

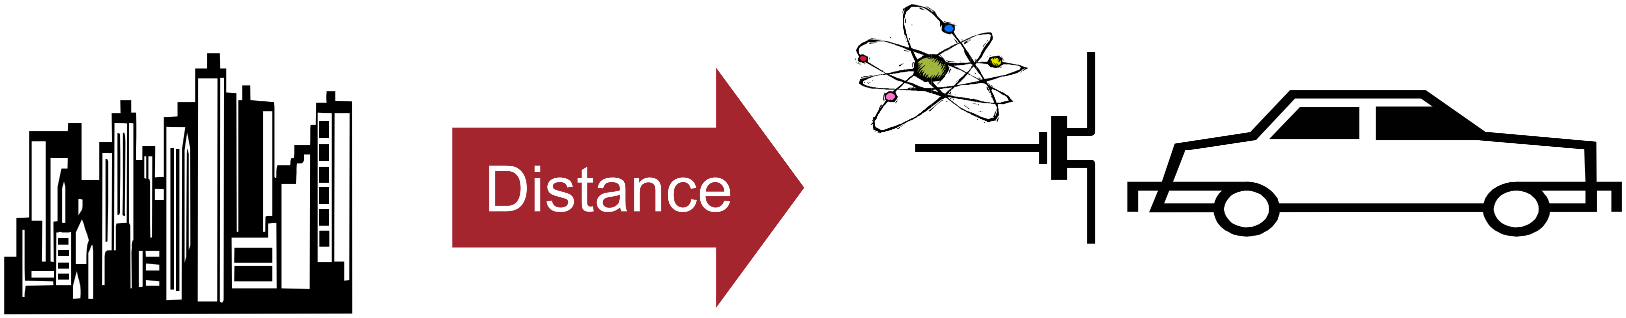

Driving away from FM transmitter→less signal
Noise from electrons → no change

Increasing numbers of gates→less signal power
Noise from electrons → no change

Sandia National Laboratories

Scientific Supercomputer Limits

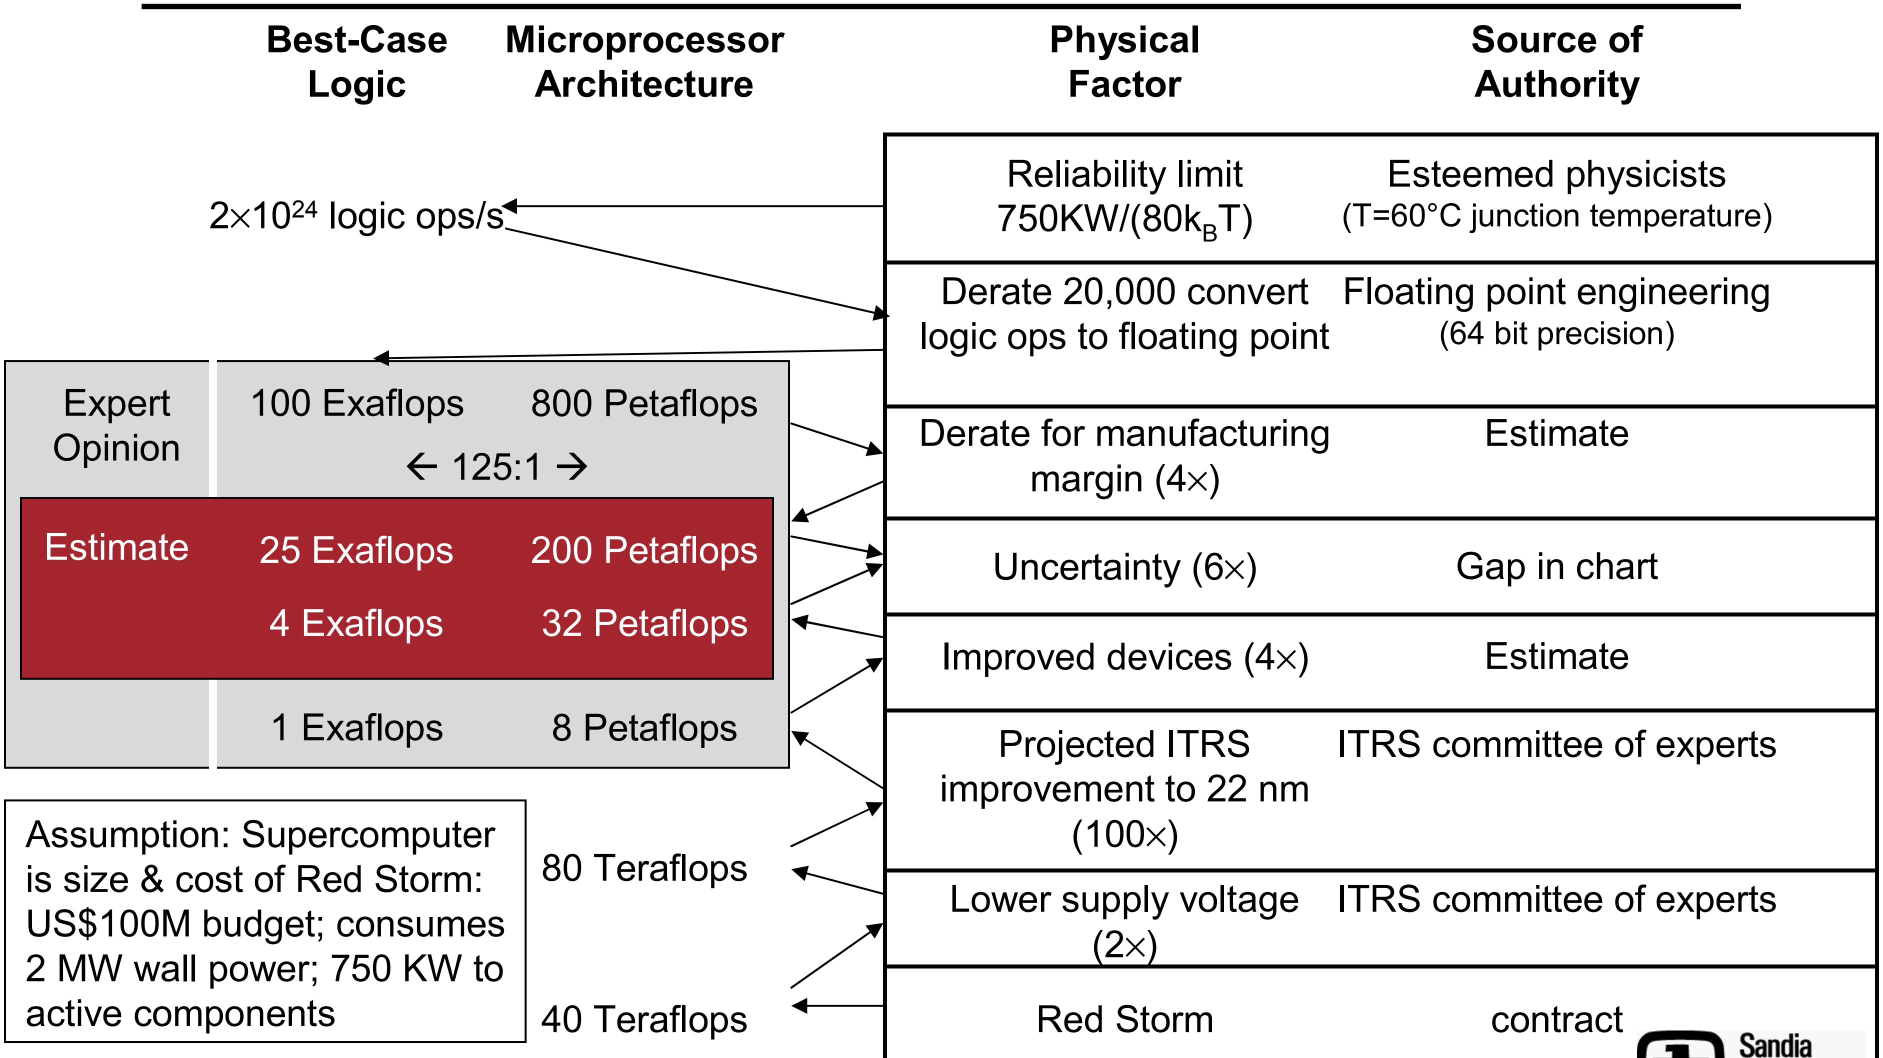

Sandia National Laboratories

Outline

- An Exemplary Zettaflops Problem
- The Limits of Moore's Law
- Beyond Moore's Law
  - Industry's Plans
  - Nanotech and Reversible Logic
- Conclusions

Sandia National Laboratories

Proceeding

- So industry has plans to extend Moore's Law, right?
  - Next slide shows ITRS Emerging Research Devices (ERD), the devices under consideration by industry
  - All are either hotter, bigger, or slower
  - Erik is now on ITRS ERD committee
- What is scientifically feasible for Gov't funding?
  - Nanotechnology
    - Efforts all over
  - Reversible logic
    - Odd name for a method of cutting power below k_BT
    - Not currently embraced by industry

Sandia National Laboratories

ITRS Device Review 2016

Table 65 Estimated Parameters for Emerging Research Devices and Technologies in the year 2016

Technology	T_{min} sec	T_{max} sec	CD_{min} m	CD_{max} m	Energy J/op	Cost min $/gate	Cost max $/gate
Si CMOS	3E-11[121]	1E-6	8E-9	5E-6	4E-18	1E-11	3E-3
RSFQ	1E-12	5E-11	3E-7	1E-6	2E-18	1E-3	1E-2
Molecular	1E-8	1E-3	1E-9	5E-9	1E-20	1E-12	1E-10
Plastic	1E-4	1E-3	1E-4	1E-3	4E-18	1E-7	1E-6
Optical (digital, all optical)	1E-16	1E-12	2E-7	2E-6	1E-12	1E-3	1E-2
NEMS (conservative)	1E-7	1E-3	1E-8	1E-7	1E-21	1E-8[122]	1E-5
Biologically Inspired	1E-13	1E-4	6E-6	5E-5	3E-25	5E-4	3E-1
Quantum	1E-16	1E-15	1E-8	1E-7	1E-21	1E3	1E5

In this table T stands for system cycle time (switching time), CD stands for critical dimension (e.g., physical gate length), Energy is the intrinsic operational energy of one device, and Cost is defined as $ per gate.

Shaded = ↑ Slower ↑ ↑ Larger ↑ Hotter

Sandia National Laboratories

Outline

- An Exemplary Zettaflops Problem
- The Limits of Moore's Law
- Beyond Moore's Law
  - Industry's Plans
  - Nanotech and Reversible Logic
- Conclusions

An Exemplary Device: Quantum Dots

- Pairs of molecules create a memory cell or a logic gate

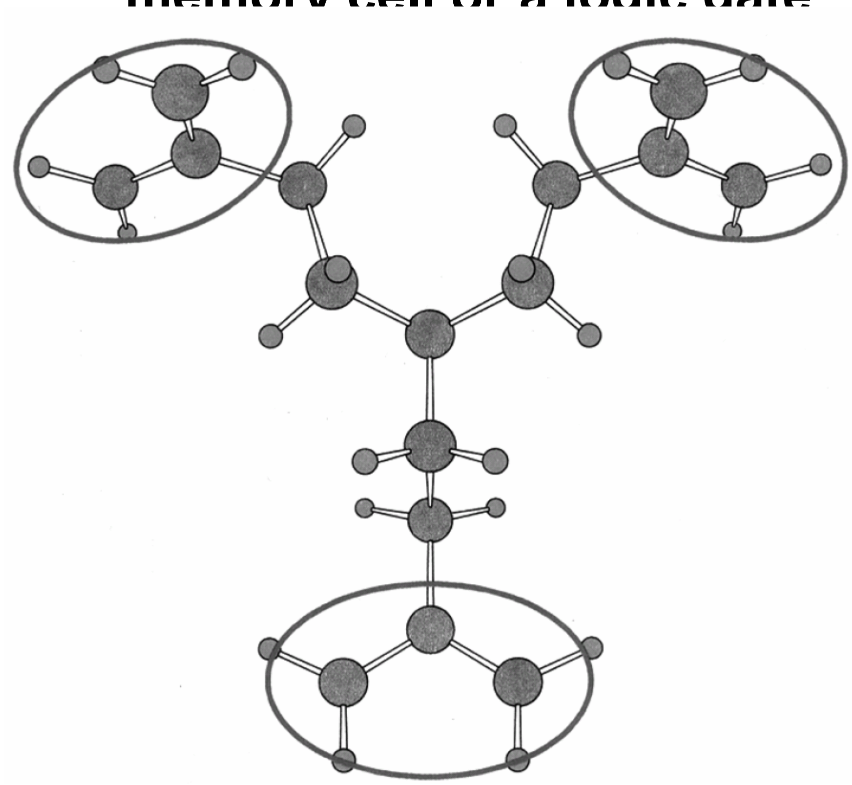

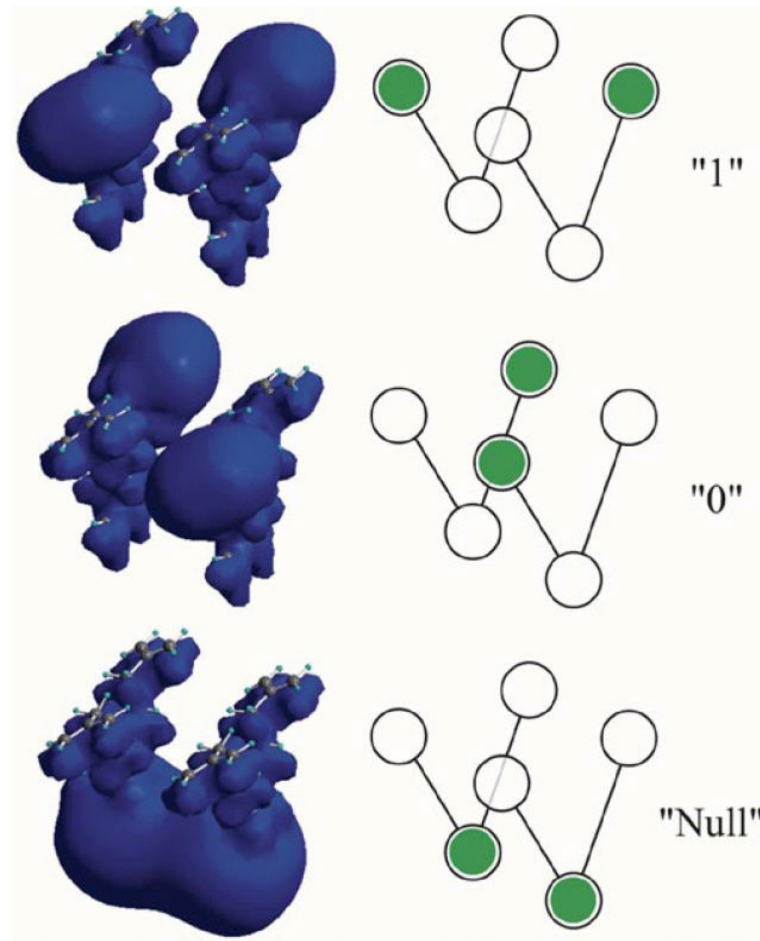

Ref. “Clocked Molecular Quantum-Dot Cellular Automata,” Craig S. Lent and Beth Isaksen
IEEE TRANSACTIONS ON ELECTRON DEVICES, VOL. 50, NO. 9, SEPTEMBER 2003

phere Simulation at a Zettaflops

Supercomputer is 211K chips, each with 70.7K nodes of 5.77K cells of 240 bytes; solves 86T=44.1Kx44.1Kx 44.1K cell problem.
System dissipates 332KW from the faces of a cube 1.53m on a side, for a power density of 47.3KW/m^2. Power: 332KW active components; 1.33MW refrigeration; 3.32MW wall power; 6.65MW from power company.
System has been inflated by 2.57 over minimum size to provide enough surface area to avoid overheating.
Chips are at 99.22% full, comprised of 7.07G logic, 101M memory decoder, and 6.44T memory transistors.
Gate cell edge is 34.4nm (logic) 34.4nm (decoder); memory cell edge is 4.5nm (memory).
Compute power is 768 EFLOPS, completing an iteration in 224µs and a run in 9.88s.

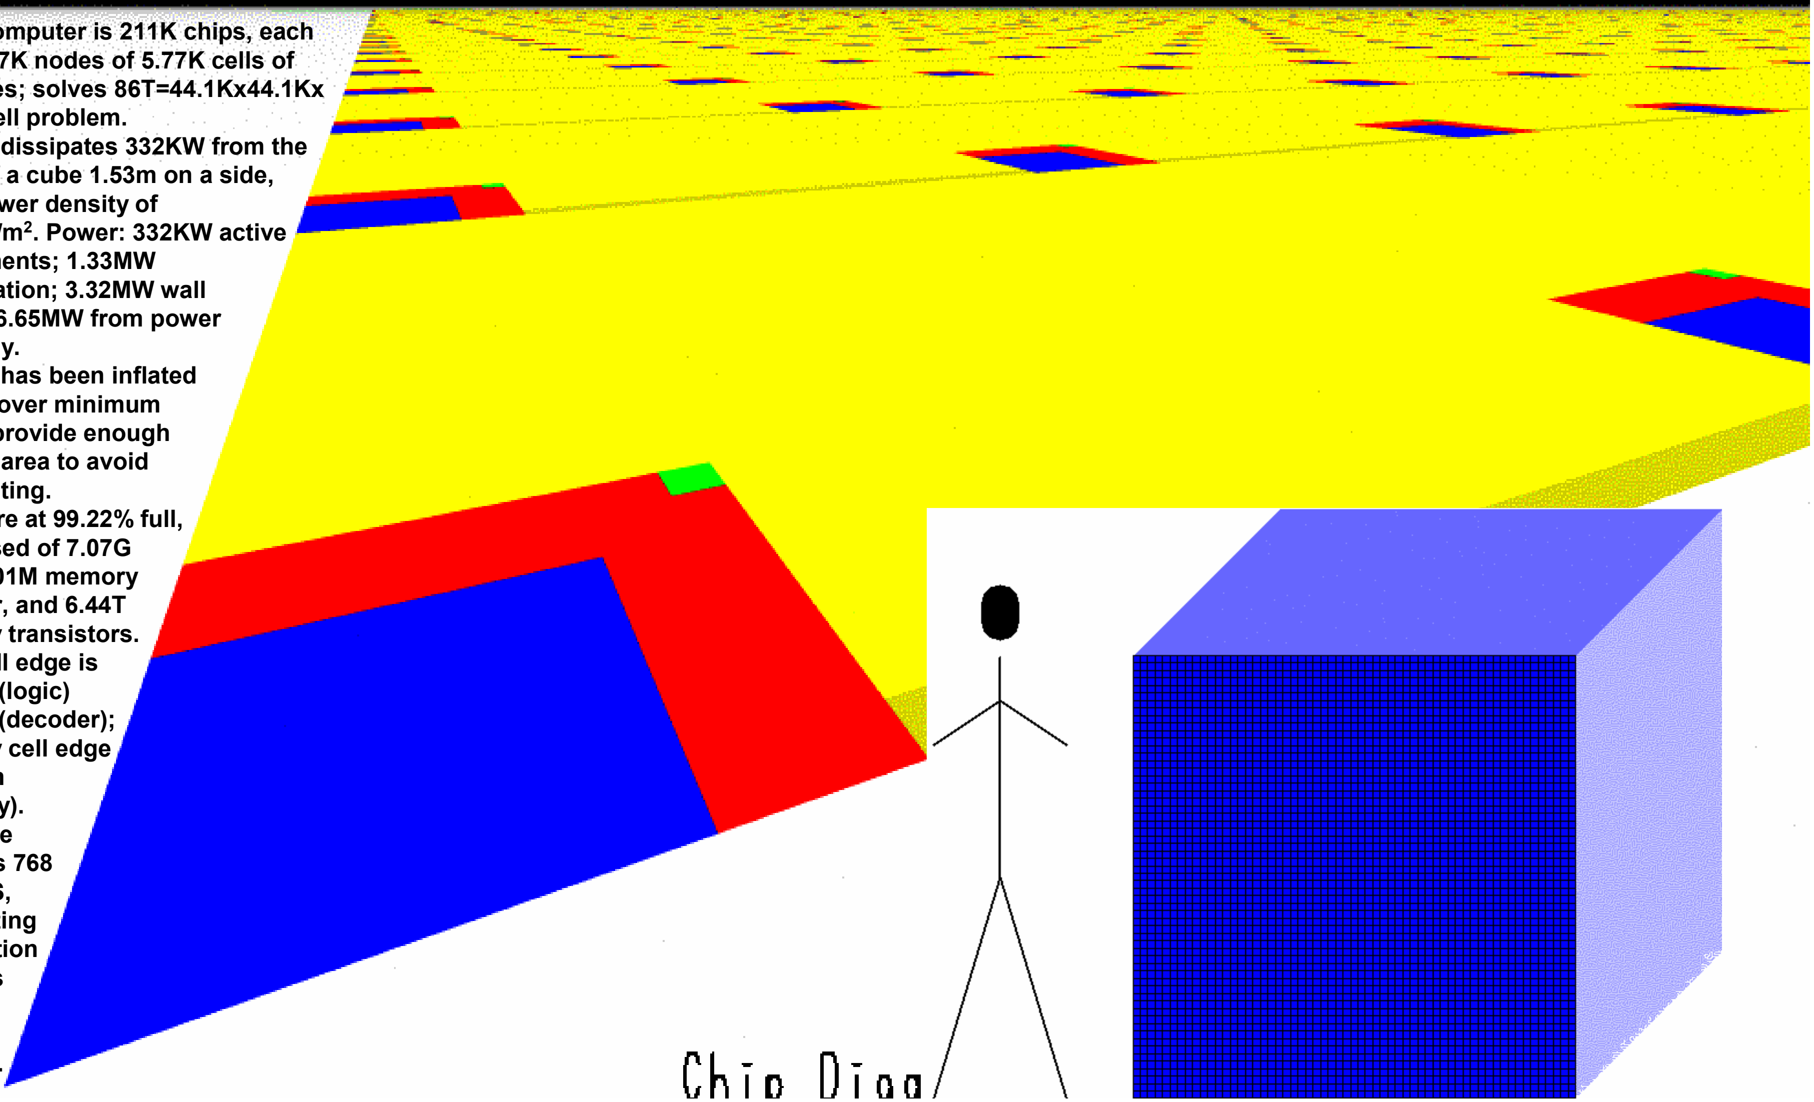

Outline

- An Exemplary Zettaflops Problem
- The Limits of Moore's Law
- Beyond Moore's Law
  - Industry's Plans
  - Nanotech and Reversible Logic
- Conclusions

Conclusions

- Important applications exist to 1 Zettaflops
- Performance of $100M μP-based supercomputers will rise to only ~100 Petaflops
  - This would meet current ASC demand
  - Will take decades to reach this level
  - But then again, 100 Teraflops was supposed to as well
- Advanced Architectures (e. g. PIM) will rise to ~10 Exaflops
  - Sandia has proposal outstanding to deploy Cyclops PIM-based system
- Nanotech and Reversible logic good to perhaps 1 Zettaflops