Review Approval

Prepare Request
Search Requests
Generate Reports
Approvals
Help
Wizard

Search Requests
New Search
Refine Search
Search Results

Clone Request
Edit Request
Cancel Request

Search Detail

Submittal Details

Document Info

Title : Petaflops Exaflops and Zettaflops for Climate Modeling
Document Number : 5240487 SAND Number : 2006-1126 C
Review Type : Electronic Status : Approved
Sandia Contact : DEBENEDICTIS,ERIK P. Submittal Type : Conference Paper
Requestor : DEBENEDICTIS,ERIK P. Submit Date : 02/24/2006
Comments : Invited presentation. Full title of workshop is "8th International Workshop on Next Generation Climate Models for Advanced High Performance Computing Facilities held jointly with the 3rd International Workshop on the Kyosei Project"
Peer Reviewed? : N

Author(s)

DEBENEDICTIS,ERIK P.

Event (Conference/Journal/Book) Info

Name : 8th International Workshop on Next Generation Climate Models for Advanced High Performance Computin
City : Albuquerque State : NM Country : USA
Start Date : 02/23/2006 End Date : 02/25/2006

Partnership Info

Partnership Involved : No
Partner Approval : Agreement Number :

Patent Info

Scientific or Technical in Content : Yes
Technical Advance : No TA Form Filed : No
SD Number :

Classification and Sensitivity Info

Title : Unclassified-Unlimited Abstract : Document : Unclassified-Unlimited
Additional Limited Release Info : None.
DUSA : DIS-CS

Routing Details

Role	Routed To	Approved By	Approval Date
Manager Approver	PUNDIT,NEIL D.	PUNDIT,NEIL D.	02/24/2006
Conditions:			
Administrator Approver	LUCERO,ARLENE M.	CANDELARIA,ALYCIA D.	02/21/2007

SAND 2006-1126C

Petaflops, Exaflops, and Zettaflops for Climate Modeling

The 8th International Workshop on Next Generation Climate Models for Advanced High Performance Computing Facilities

Erik P. DeBenedictis

Sandia National Laboratories

February 23, 2006

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.

Sandia National Laboratories

Applications and $100M Supercomputers

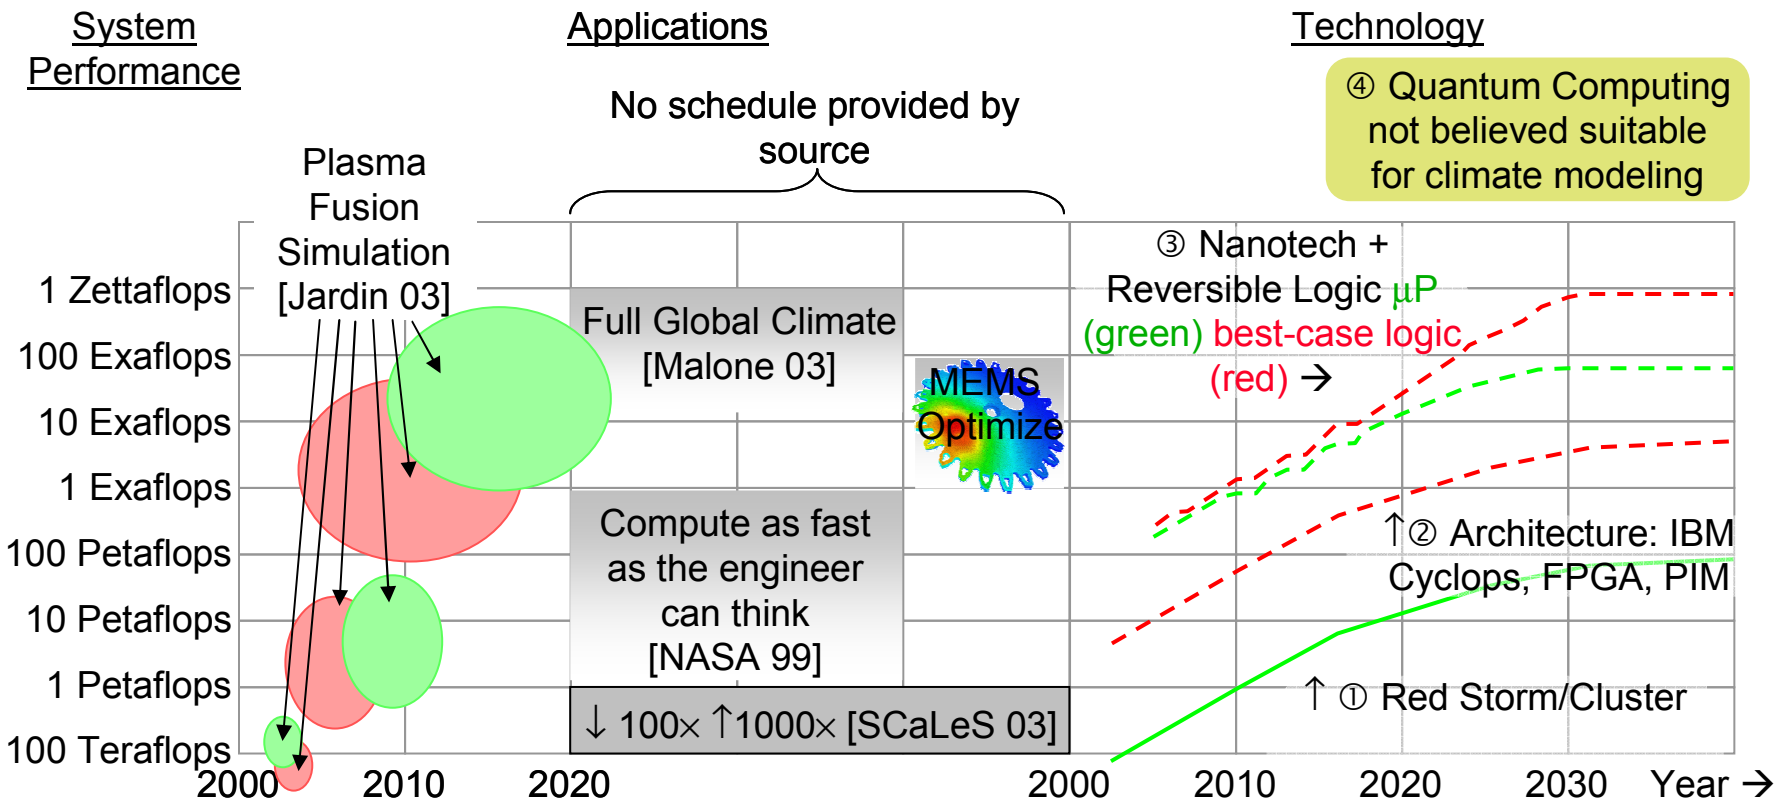

[Jardin 03] S.C. Jardin, "Plasma Science Contribution to the SCaLeS Report," Princeton Plasma Physics Laboratory, PPPL-3879 UC-70, available on Internet.
[Malone 03] Robert C. Malone, John B. Drake, Philip W. Jones, Douglas A. Rotman, "High-End Computing in Climate Modeling," contribution to SCaLeS report.
[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

- **Exemplary Zettaflops Problems**
- **The Limits of Moore's Law**
  - **Computer Physics for Climate Scientists**
- **New Computing Paradigms**
  - **Reversible Logic and Nanotech**
- **Feasibility of Climate Modeling at 1 Zettaflops**
- **Conclusions**

Sandia National Laboratories

Simulation of Global Climate

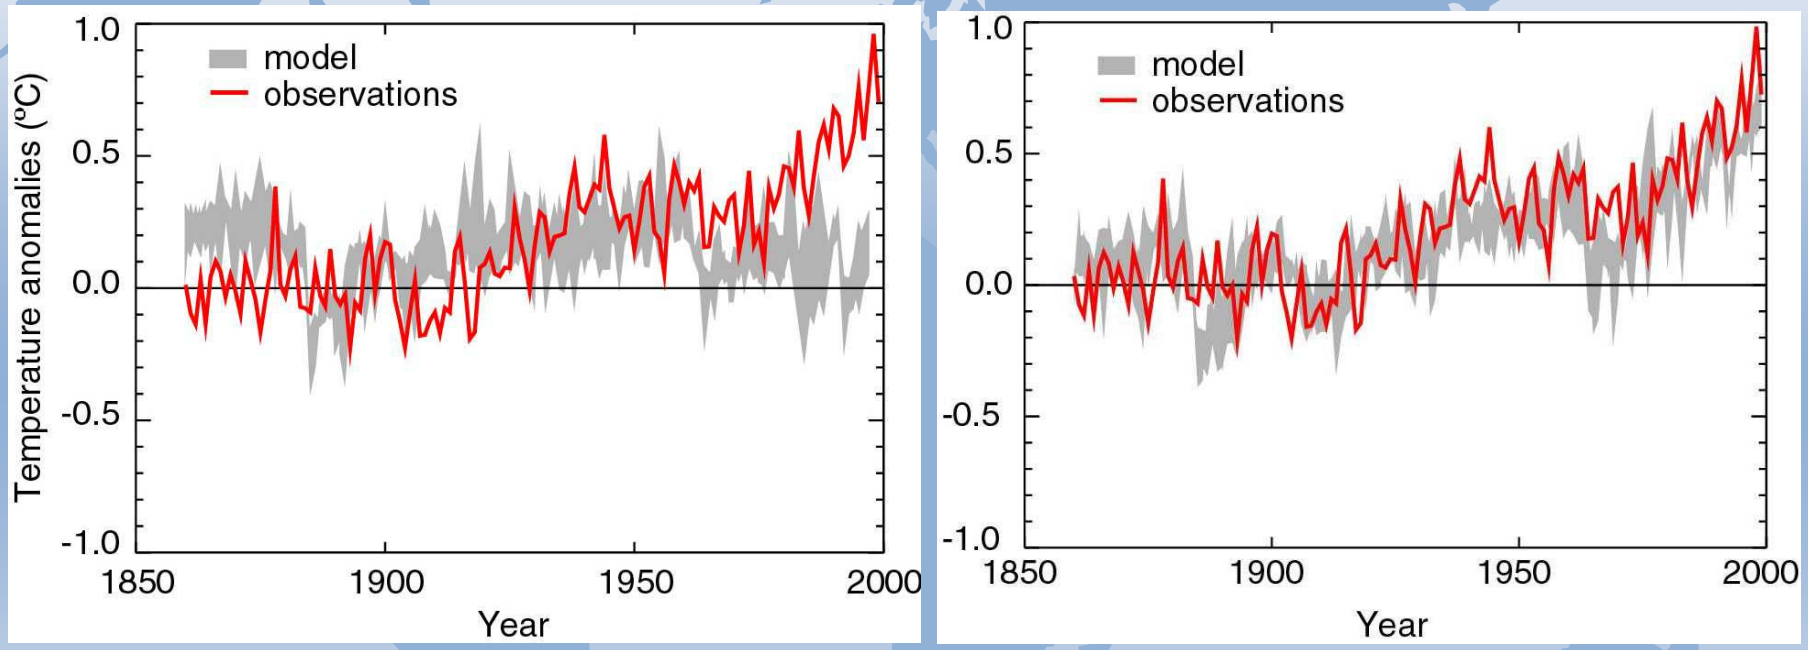

Stott et al, Science 2000

"Simulations of the response to natural forcings alone … do not explain the warming in the second half of the century"

"..model estimates that take into account both greenhouse gases and sulphate aerosols are consistent with observations over this*period" - IPCC 2001

Office of Science U.S. DEPARTMENT OF ENERGY

SciDAC Scientific Discovery through Advanced Computing

The World's Greatest Science Protecting America

Los Alamos NATIONAL LABORATORY EST.1943

FLOPS Increases for Global Climate

	Issue	Scaling
1 Zettaflops	Ensembles, scenarios 10×	Embarrassingly Parallel
100 Exaflops	Run length 100×	Longer Running Time
1 Exaflops	New parameterizations 100×	More Complex Physics
10 Petaflops	Model Completeness 100×	More Complex Physics
100 Teraflops	Spatial Resolution $10^4\times$ ($10^3\times$-$10^5\times$)	Resolution
10 Gigaflops	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

Outline

- Exemplary Zettaflops Problems
- The Limits of Moore's Law
  - Computer Physics for Climate Scientists
- New Computing Paradigms
  - Reversible Logic and Nanotech
- Feasibility of Climate Modeling at 1 Zettaflops
- Conclusions

Sandia National Laboratories

*** This is a Preview ***

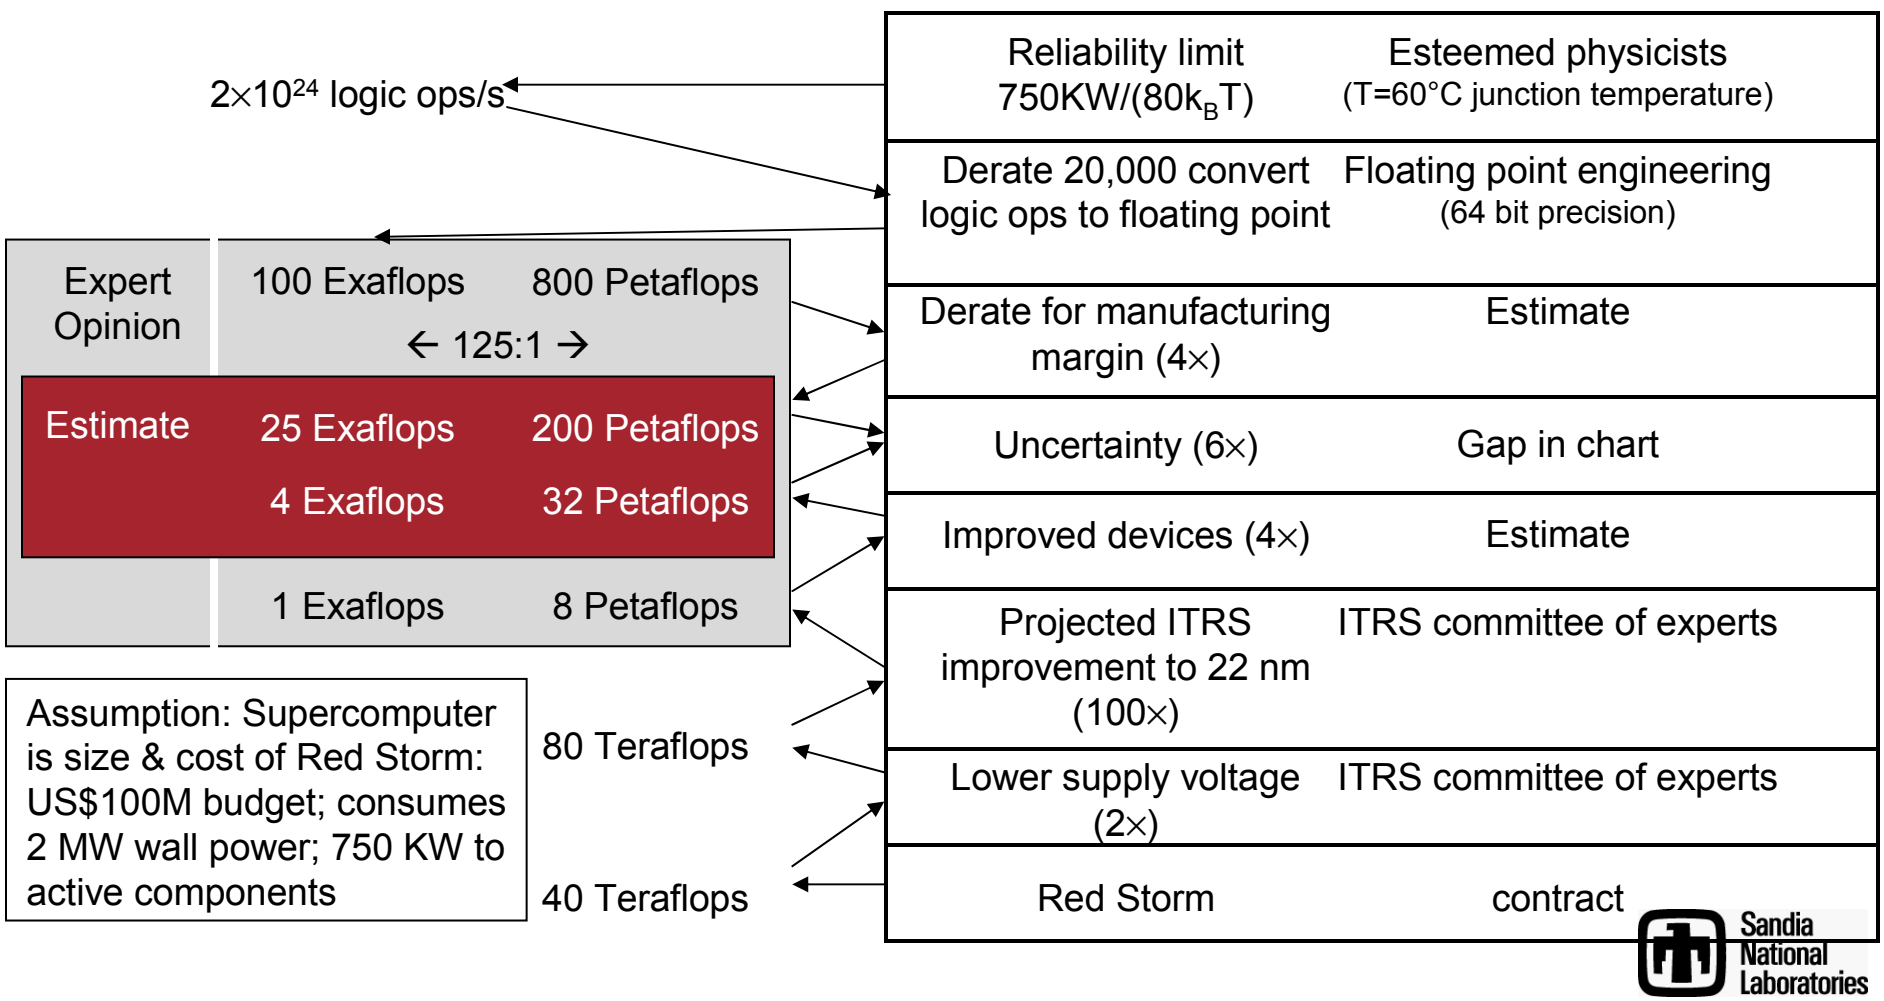

Sandia National Laboratories

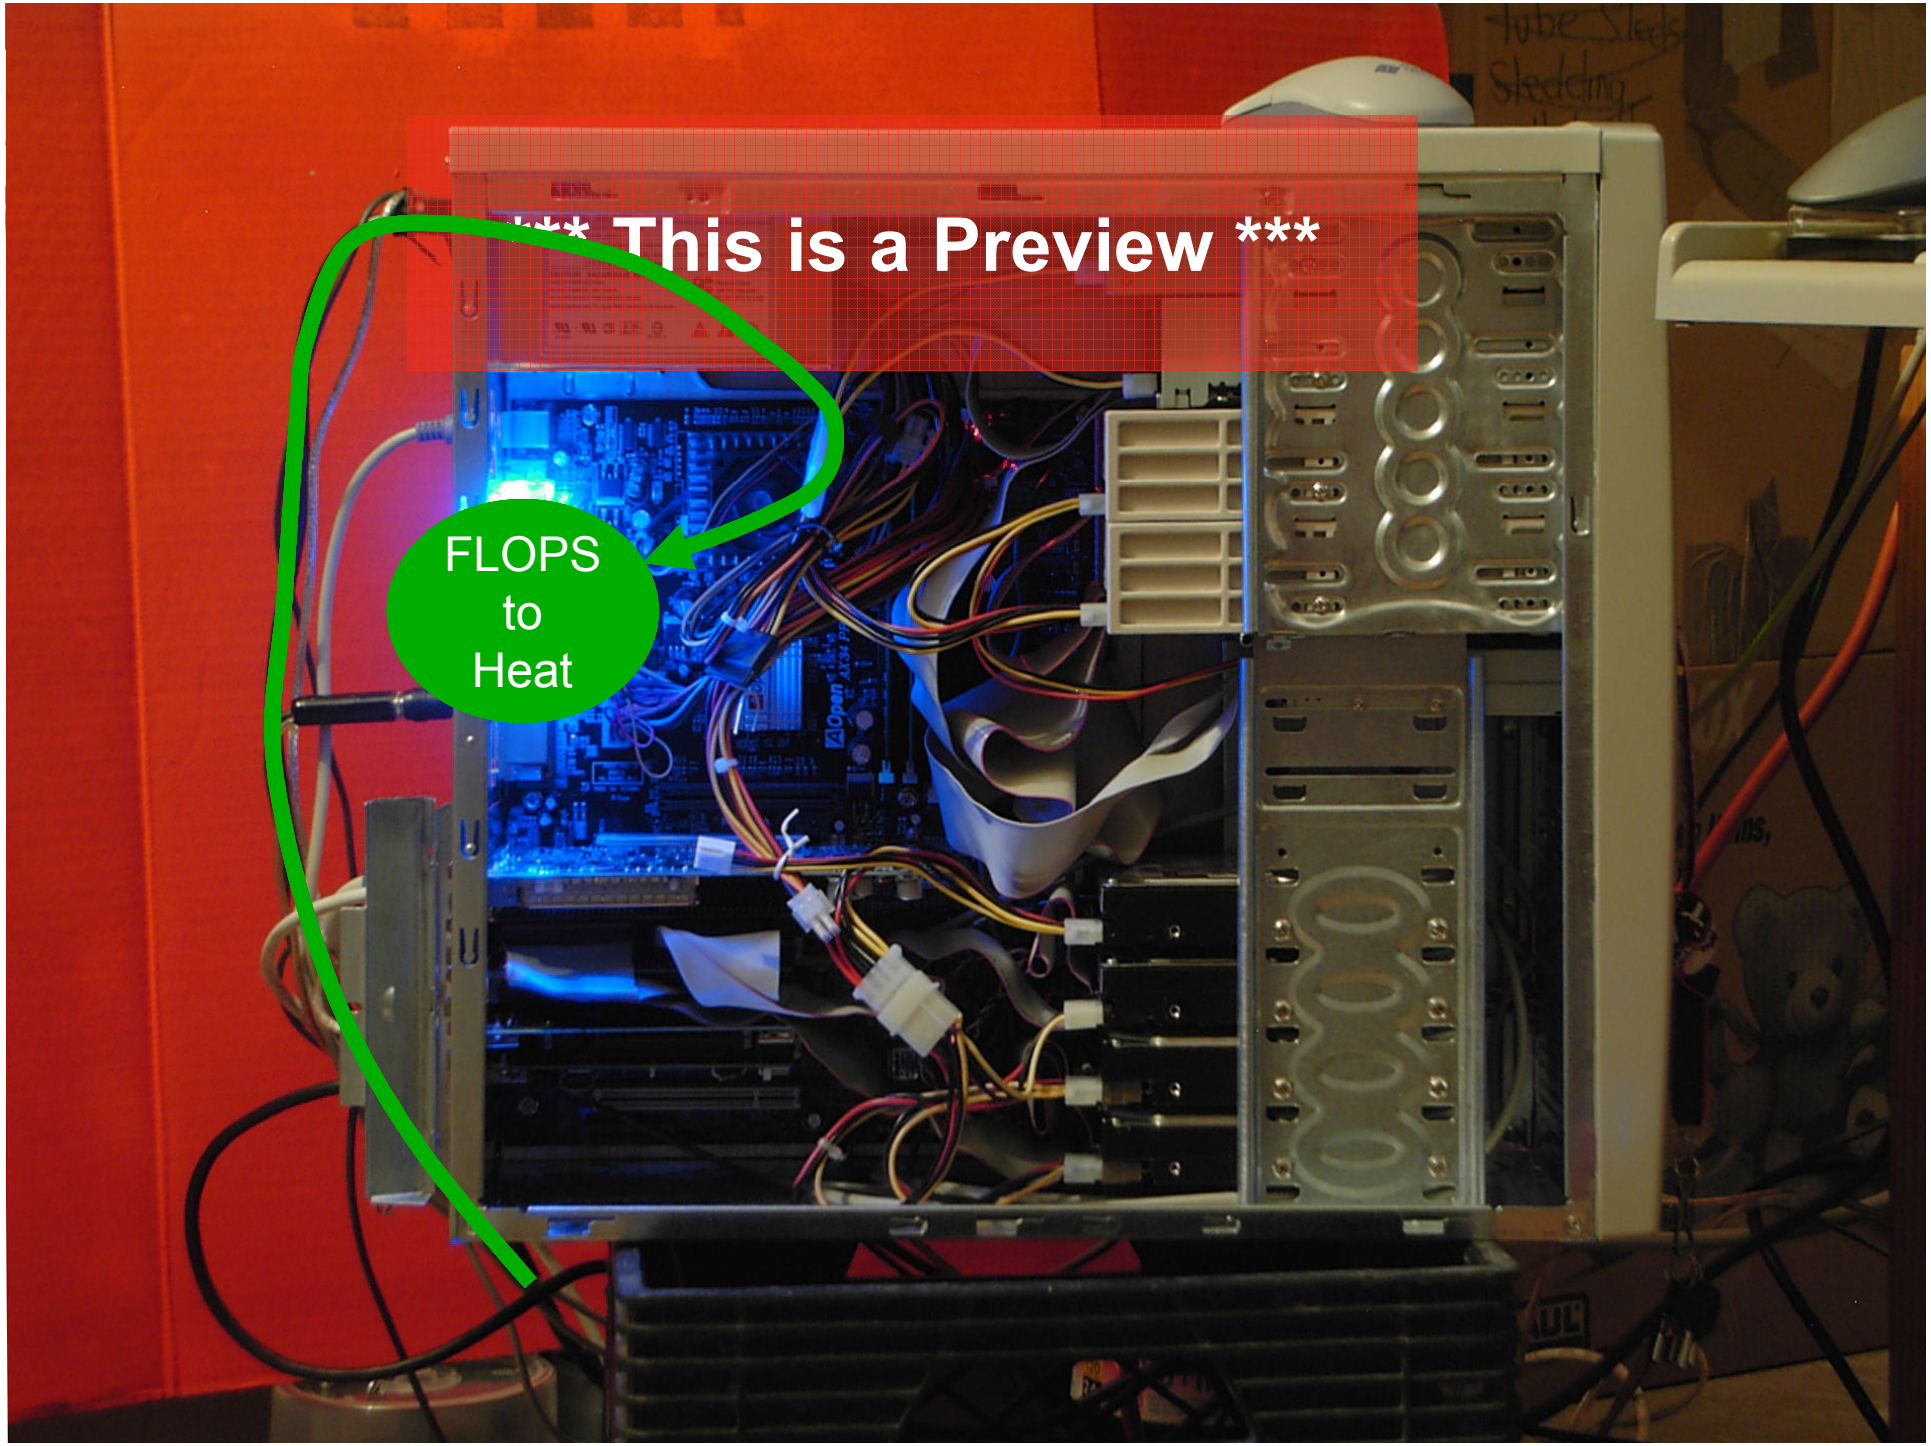
*** This is a Preview ***
FLOPS
to
Heat

Speed and Power

- If we have parallel algorithms…
  - Power per floating point operation is first order effect, because entire physical structure of machine is proportional to power
  - Speed is a second order effect because it can be traded for parallelism
- Transistor roadmap →

International Technology Roadmap for Semiconductors (ITRS), see http://public.itrs.net

Sandia National Laboratories

Semiconductor Roadmap

YEAR OF PRODUCTION	2010	2013	2016
DRAM ½ PITCH (nm)	45	32	22
MPU / ASIC ½ PITCH (nm)	50	35	25
MPU PRINTED GATE LENGTH (nm)	25	18	13
MPU PHYSICAL GATE LENGTH (nm)	18	13	9
Physical gate length high-performance (HP) (nm) [1]	18	13	9
Equivalent physical oxide thickness for high-performance T_{ox} (EOT)(nm) [2]	0.5-0.8	0.4-0.6	0.4-0.5
Gate depletion and quantum effects electrical thickness adjustment factor (nm) [3]	0.5	0.5	0.5
T_{ox} electrical equivalent (nm) [4]	1.2	1.0	0.9
Nominal power supply voltage (V_{dd}) (V) [5]	0.6	0.5	0.4
Nominal high-performance NMOS sub threshold leakage current, I_{sdleak} (at 25°C) (μA/μm) [6]	3	7	10
Nominal high-performance NMOS saturation drive current, I_{dd} (at V_{dd} at 25°C) (μA/μm) [7]	1200	1500	1500
Required percent current-drive "mobility/transconductance improvement" [8]	30%	70%	100%
Parasitic source/drain resistance (Rsd) (ohm-	110	90	80
Parasitic source/drain resistance (Rsd) pe	25%	30%	35%
Parasitic capacitance percent of ideal gat	31%	36%	42%
High-performance NMOS device τ (C_{gate} * V_{dd} / I_{dd}-NMOS)(ps) [12]	0.39	0.22	0.15
Relative device performance [13]	4.5	7.2	10.7
Energy per (W/L_{gate}=3) device switching transition (C_{gate}*(3*L_{gate})*V^2) (fJ/Device) [14]	0.015	0.007	0.002
Static power dissipation per (W/Lgate=3) device (Watts/Device) [15]	9.7E-08	1.4E-07	1.1E-07

1,000 k_BT/transistor

White—Manufacturable Solutions Exist, and Are Being Optimized

Yellow—Manufacturable Solutions are Known

Red—Manufacturable Solutions are NOT Known

Sandia National Laboratories

Computer Physics for Climate Scientists

- The 0's and 1's in a computer must be characterized by enough energy to be reliably distinguished from thermal noise
- Random motion is $k_B T$ per gate time
- To be reliable, 0's and 1's must have 70-100 times $k_B T$
  - Error p = $e^{-70} - e^{-100}$

- Imagine hypothetical cloud droplets, defined such that addition of N $k_B T$ causes them to evaporate
- While droplets are matter, their state as droplet vs. air and water vapor is information
- N could be as low as 70-100
- Let 0's be gray droplets
- Let 1's be blue droplets

Sandia National Laboratories

Ubiquitous Operation of Logic Gates

- All logic gates in use today consume the input signals and create output signals with energy from the power supply
  - Ubiquitous but not general, see later

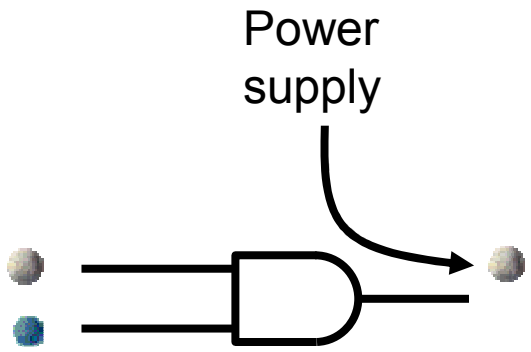

- In today's CMOS, logic signals are about $100{,}000k_BT$
- At the theoretical limit, logic signals are $70\text{-}100k_BT$
  - Note: theoretical limit applies to any logic that consumes input signals and regenerates from power supply

Sandia National Laboratories

Computer Efficiency Factors

- A floating point unit has a 52x52 multiplier array and a lot of logic to deal with exponents, exceptions, etc.
  - About 20,000 $100k_BT$ cloud droplets will be evaporated per floating point operation (arithmetic only, see →)
- A modern μP is about 1% efficient for floating point at peak
  - 99% of energy to
    - instruction decode
    - Memory bus
    - Speculative execution
  - About 2×10^6 $100k_BT$ cloud droplets per floating point operation in a μP
  - At 5% of peak, 20×10^6

Sandia National Laboratories

Physics in Source Code

- **You can figure out the limits of computer performance with source code inspection**

Reference cloud droplets:
Limit 100 k_BT
Current CMOS 100,000 k_BT

Floating Point: 20,000 cloud droplets evaporate per operation

```
      DO 100 I = 1, L
         DO 90 J = I, L
            B( I, N-L+J ) = B( I, N-L+J ) - BETA( K+I )*R( K+I, K+J )
 90      CONTINUE
100   CONTINUE
```

Microprocessor efficiency: 100× arithmetic
Peak: Divide by fraction of peak

Sandia National Laboratories

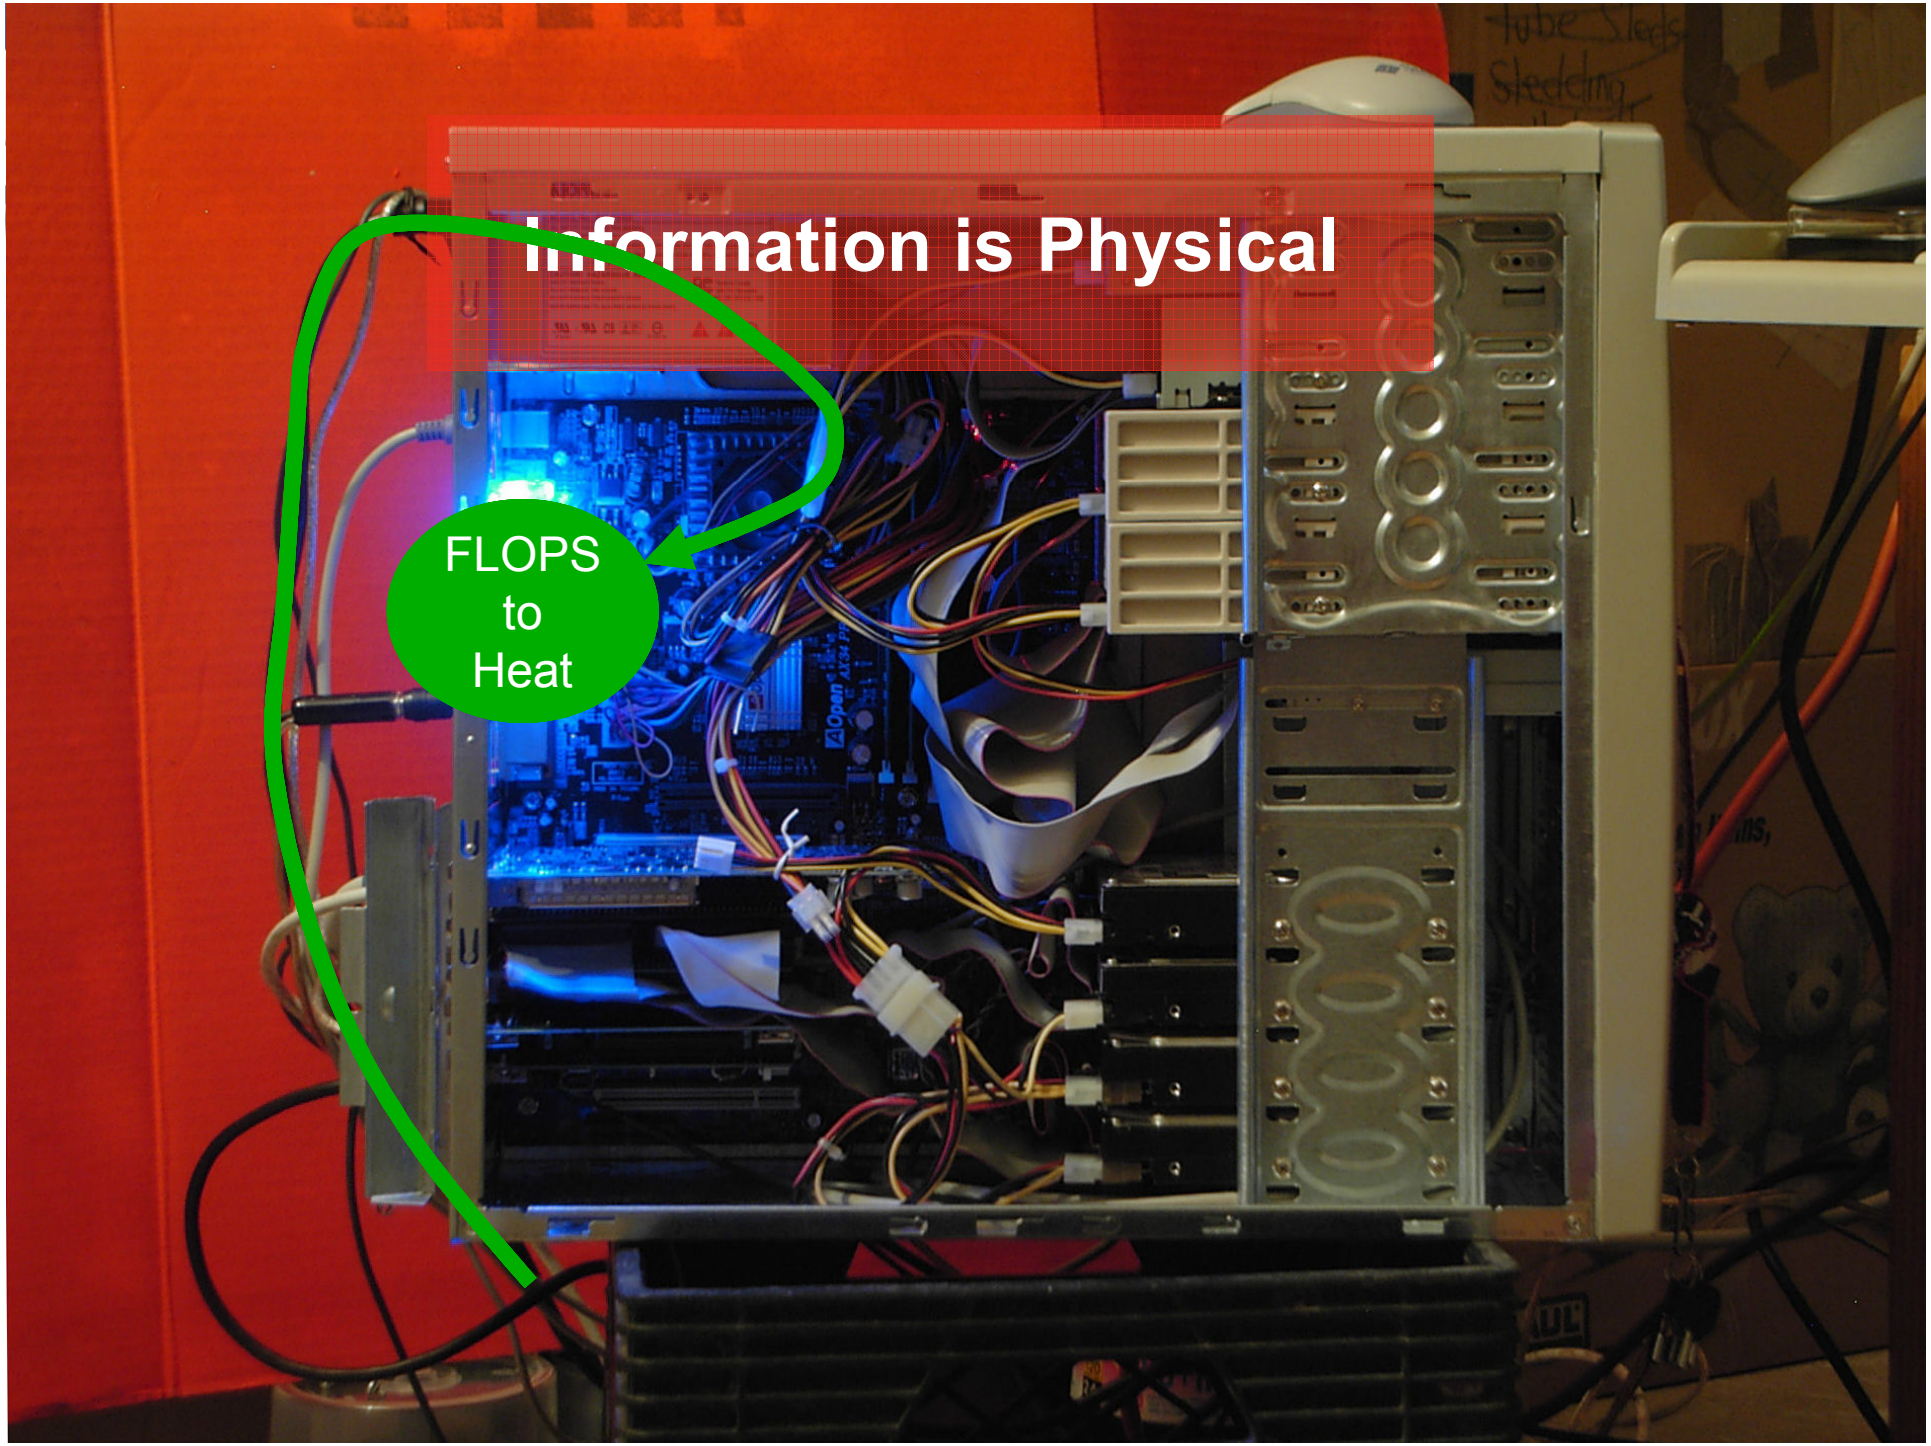
Information is Physical
FLOPS
to
Heat

Scientific Supercomputer Limits

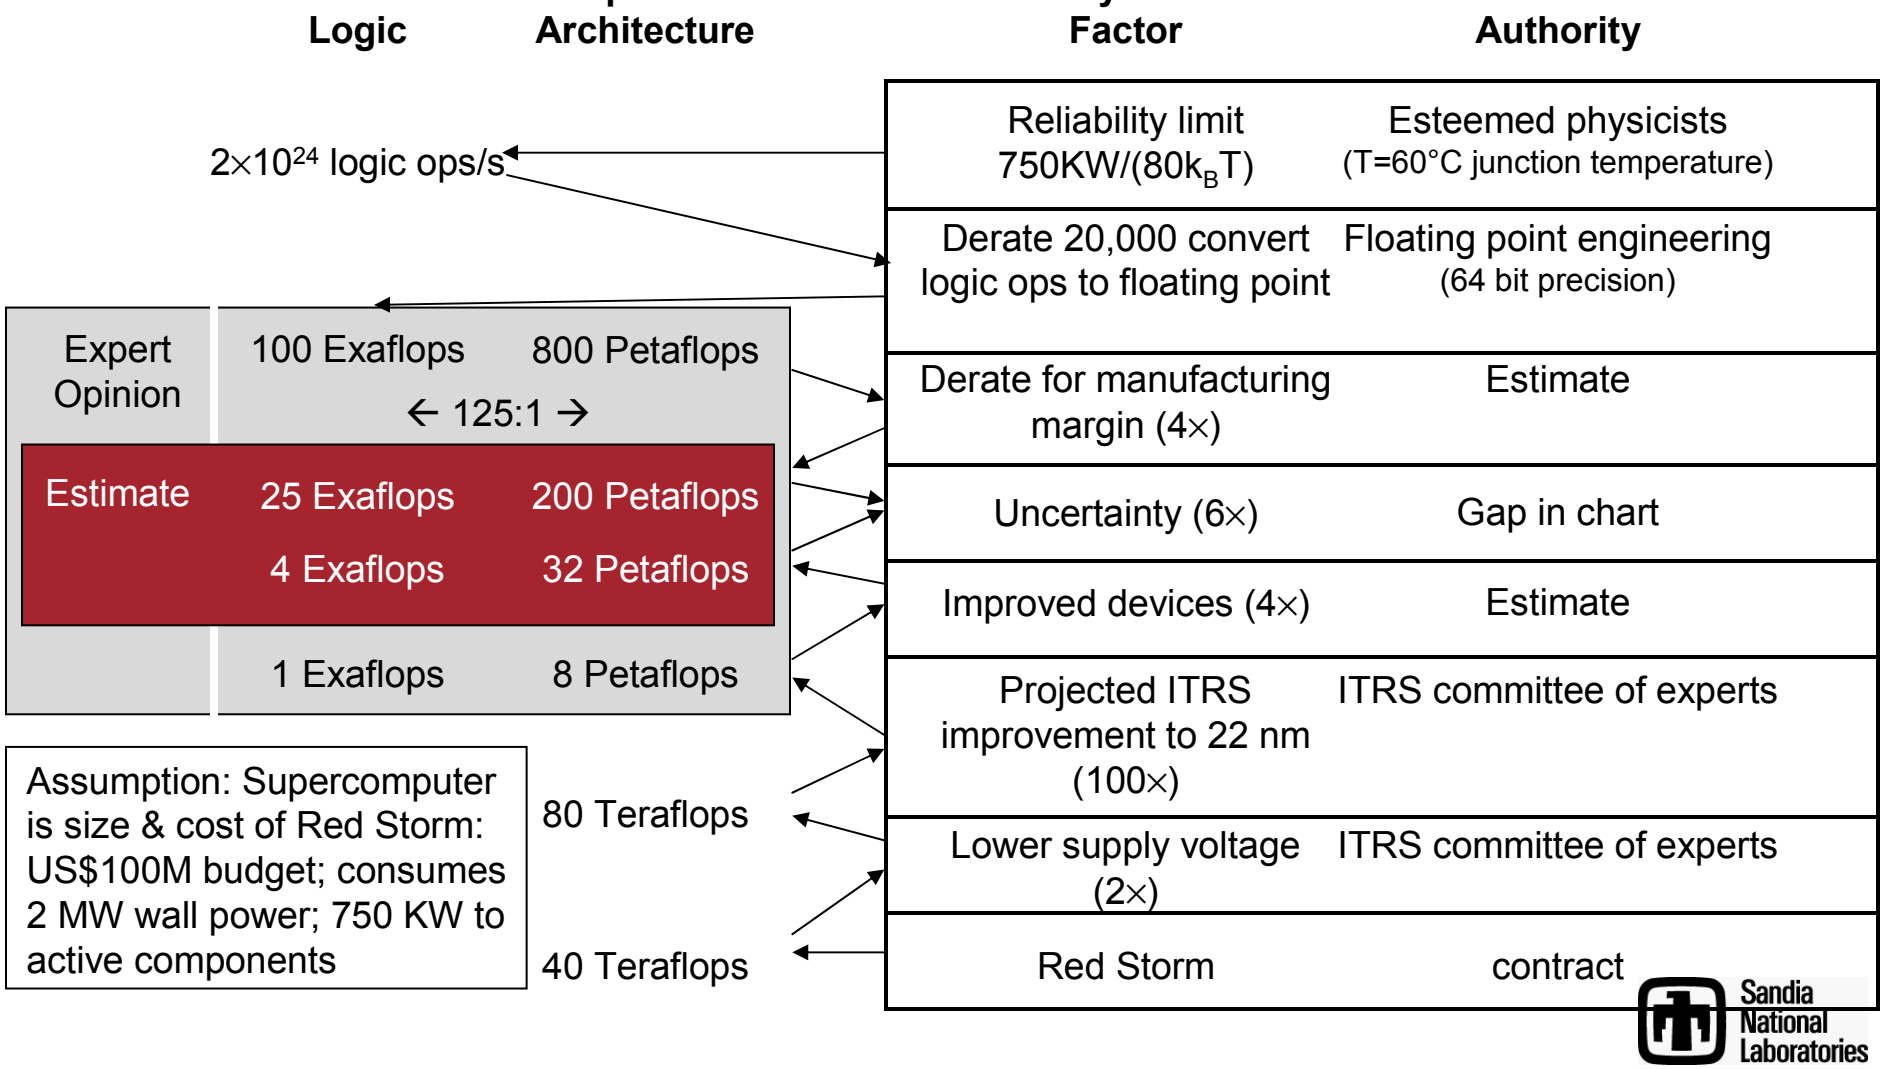

Sandia National Laboratories

Outline

- Exemplary Zettaflops Problems
- The Limits of Moore's Law
  - Computer Physics for Climate Scientists
- New Computing Paradigms
  - Reversible Logic and Nanotech
- Feasibility of Climate Modeling at 1 Zettaflops
- Conclusions

Sandia National Laboratories

A Universal Logic With No First Order Heat

- I just explained the limits of computation in a technology-independent way
- However, there is a loophole
- There are universal logic families that do not generate heat (to first order)

- Fredkin Gate
- Ctl=1, swap else no-op

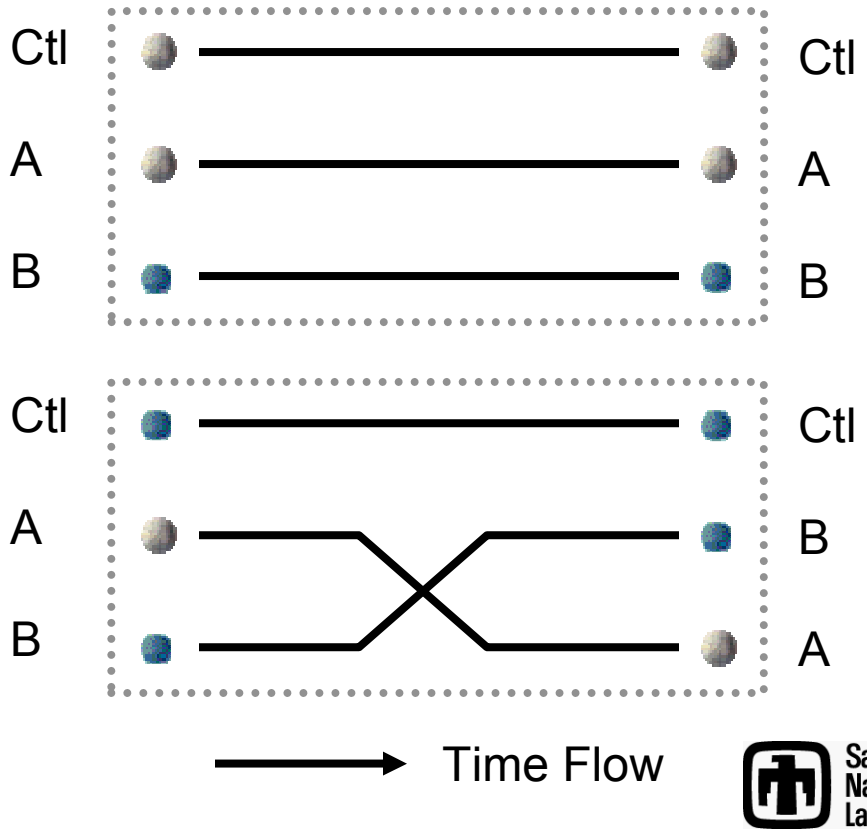

Sandia National Laboratories

Reversible Microprocessor Status

- Status
  - Subject of Ph. D. thesis
  - Chip laid out (no floating point)
  - RISC instruction set
  - C-like language
  - Compiler
  - Demonstrated on a PDE
  - However: really weird and not general to program with +=, -=, etc. rather than =

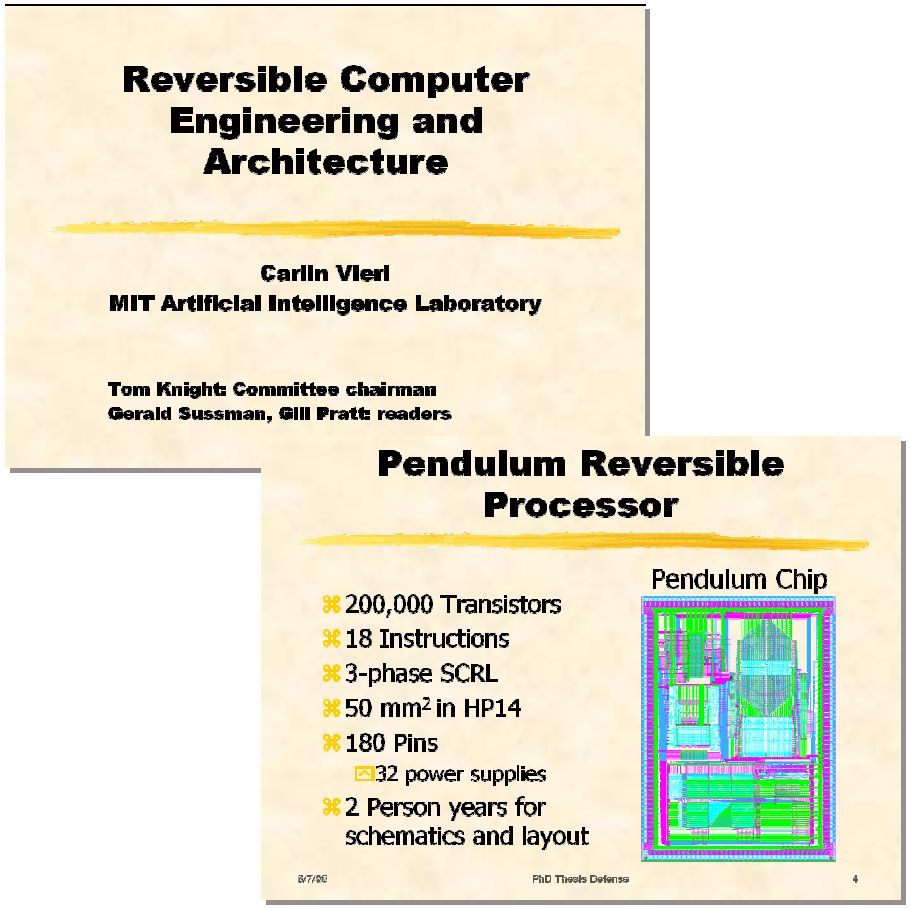

Sandia National Laboratories

Human Cost For Other Forms of Logic

- Most of us learned about logic as a teenager
- We learned about AND-OR-NOT logic, not as an instance of a general class of logics, but probably intuitively
- Computer arithmetic and programming build upon AND-OR-NOT logic
- To beat the thermodynamic limit, we need to rewind our learning to when we were a teenager and scrutinize everything computer arithmetic and programming based on a different logic basis
- This is a hard to do…in human terms

Outline

- Exemplary Zettaflops Problems
- The Limits of Moore's Law
  - Computer Physics for Climate Scientists
- New Computing Paradigms
  - Reversible Logic and Nanotech
- Feasibility of Climate Modeling at 1 Zettaflops
- Conclusions

Sandia National Laboratories

The Parallelism Issue

- **Initially, didn't meet constraints**

Scaled Climate Model

2D → 3D mesh,
one cell per processor

Parallelize cloud-resolving model and ensembles

One Barely Plausible Solution

Consider special purpose logic with fast logic and low-power memory

Consider only highest performance published nanotech device QDCA

Initial reversible nanotech

An Exemplary Device: Quantum Dots

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

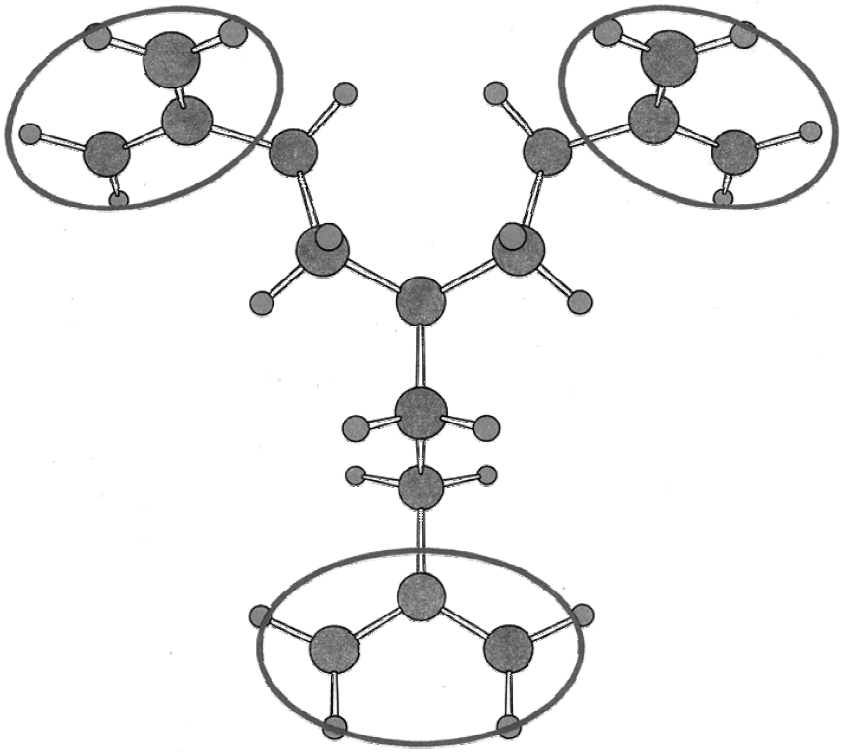

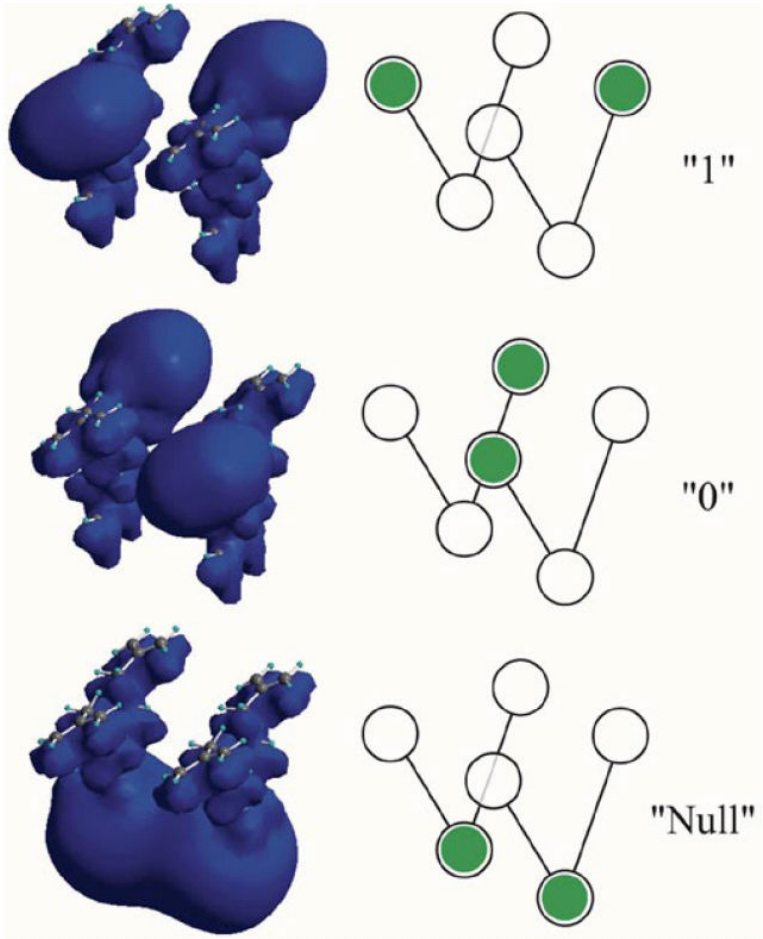

Ref. "Clocked Molecular Quantum-Dot Cellular Automata," Craig S. Lent and Beth Isaksen
IEEE TRANSACTIONS ON ELECTRON DEVICES, VOL. 50, NO. 9, SEPTEMBER 2003

Sandia National Laboratories

Performance Curve

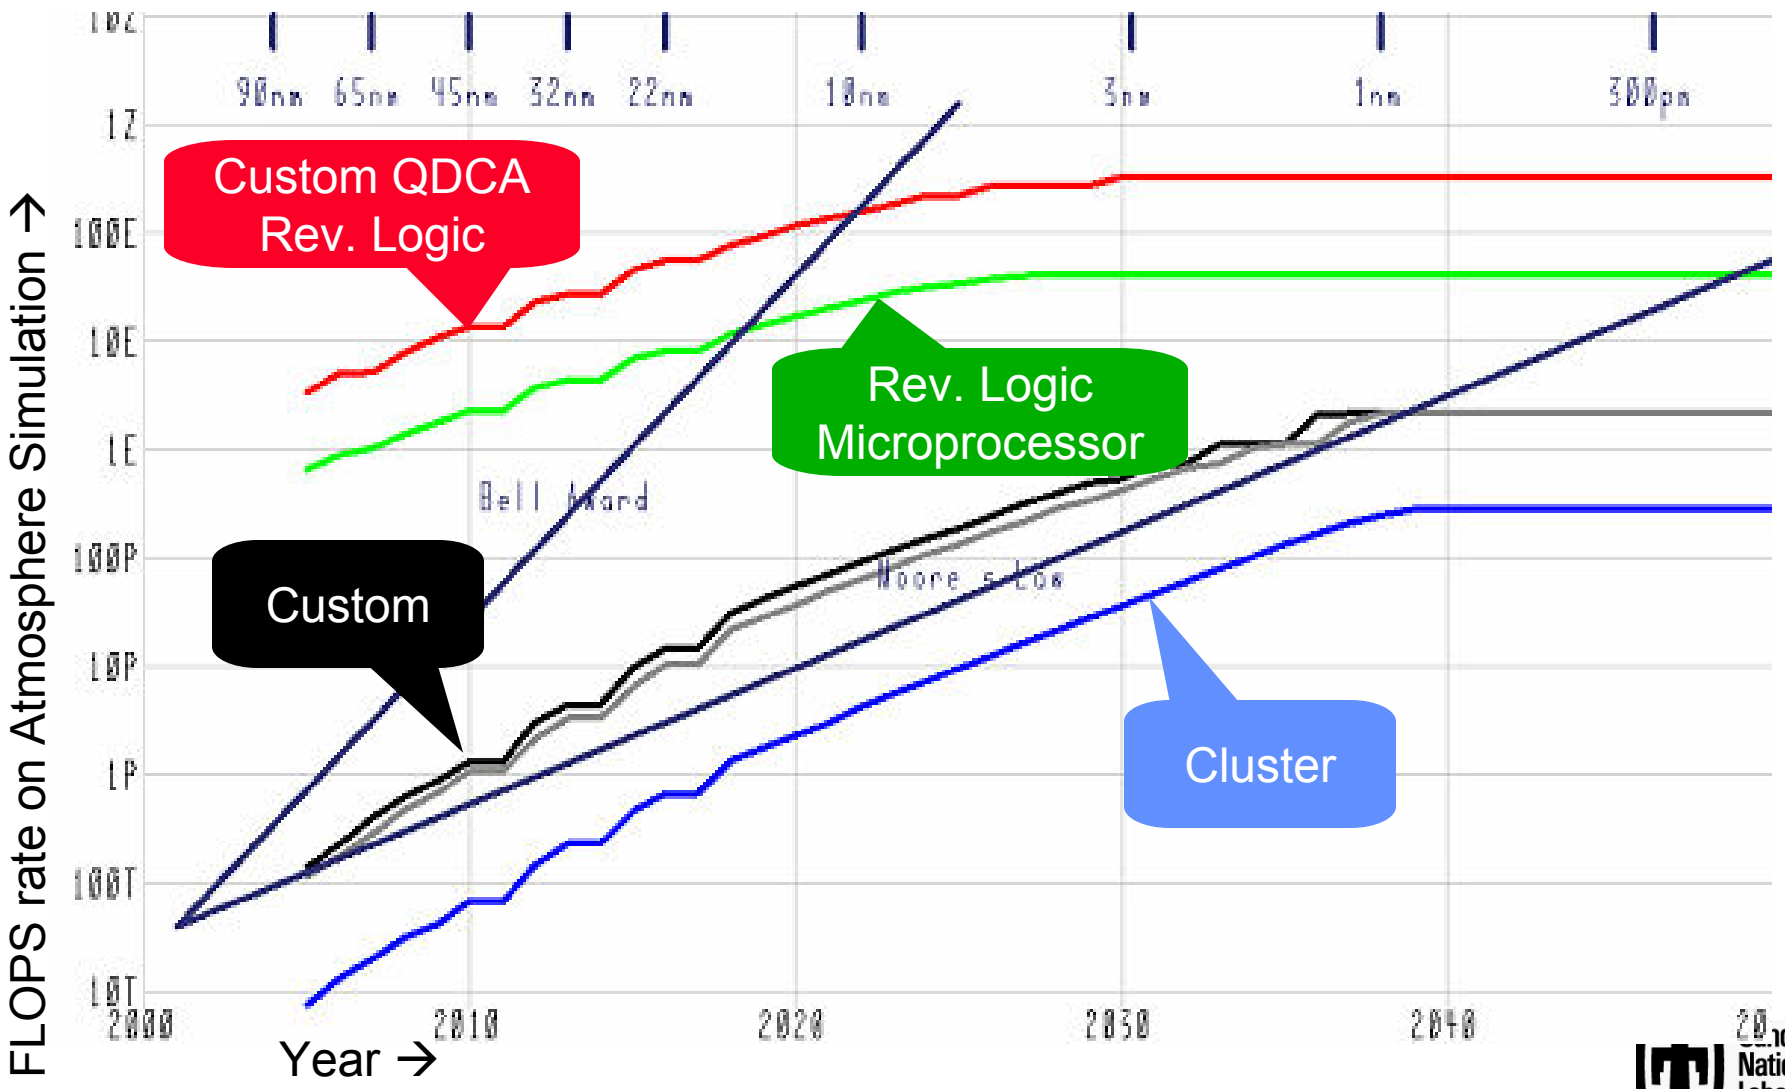

India National Laboratories

Outline

- Exemplary Zettaflops Problems
- The Limits of Moore's Law
  - Computer Physics for Climate Scientists
- New Computing Paradigms
  - Reversible Logic and Nanotech
- Feasibility of Climate Modeling at 1 Zettaflops
- Conclusions

Sandia National Laboratories

Conclusions

- Ambitious applications need up to 1 Zettaflops
- Moore's Law for devices good for another 100× performance boost
  - In power; speed less
- Architecture has 100× upside
  - But may require code tuning or rewrite
- Other models of computation exist and are under study
  - May reach 1 Zettaflops
  - May run Fortran
  - Will introduce new and uncomfortable concepts
- Conclusion: There is a rough match between computer technology and ambitious applications

Sandia National Laboratories

Outline

- Exemplary Zettaflops Problems
- The Limits of Moore's Law
  - Computer Physics for Climate Scientists
  - Industry's Plans
- New Computing Paradigms
  - Reversible Logic and Nanotech
- 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
  - Quantum computing
    - More later

Sandia National Laboratories

ITRS Device Review 2016

Technology	Speed (min-max)	Dimension (min-max)	Energy per gate-op	Comparison
CMOS	30 ps-1 μs	8 nm-5 μm	4 aJ	
RSFQ	1 ps-50 ps	300 nm- 1μm	2 aJ	Larger
Molecular	10 ns-1 ms	1 nm- 5 nm	10 zJ	Slower
Plastic	100 μs-1 ms	100 μm-1 mm	4 aJ	Larger+Slower
Optical	100 as-1 ps	200 nm-2 μm	1 pJ	Larger+Hotter
NEMS	100 ns-1 ms	10-100 nm	1 zJ	Slower+Larger
Biological	100 fs-100 μs	6-50 μm	.3 yJ	Slower+Larger
Quantum	100 as-1 fs	10-100 nm	1 zJ	Larger

Data from ITRS ERD Section.

Sandia National Laboratories

Atmosphere 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.