Sandia National Laboratories report SAND2004-5871C
Unclassified Unlimited Release

The Path To Extreme Computing

Erik P. DeBenedictis, Organizer
Sandia National Laboratories

Los Alamos Computer Science Institute Symposium 2004

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.

*** This is a Preview ***

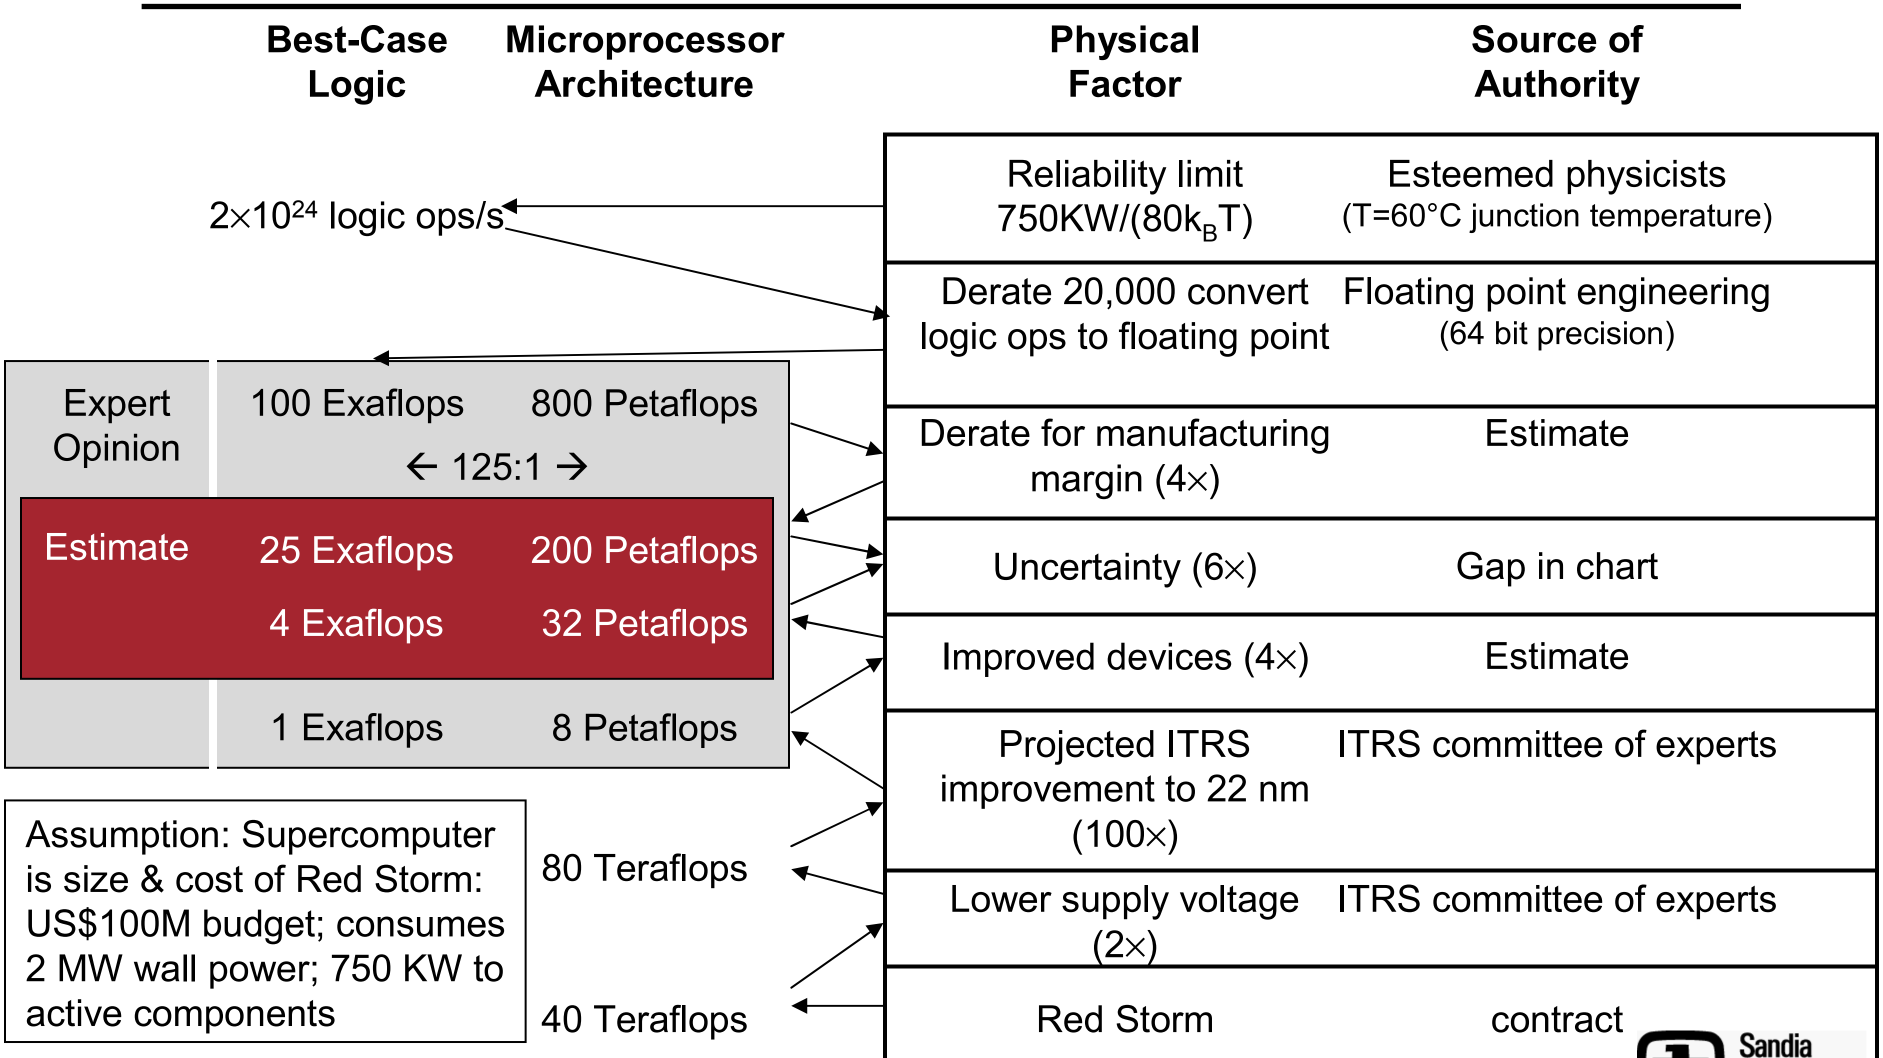

Thermal Noise Limit

This logical irreversibility is associated with physical irreversibility and requires a minimal heat generation, per machine cycle, typically of the order of kT for each irreversible function.
– R. Landauer 1961

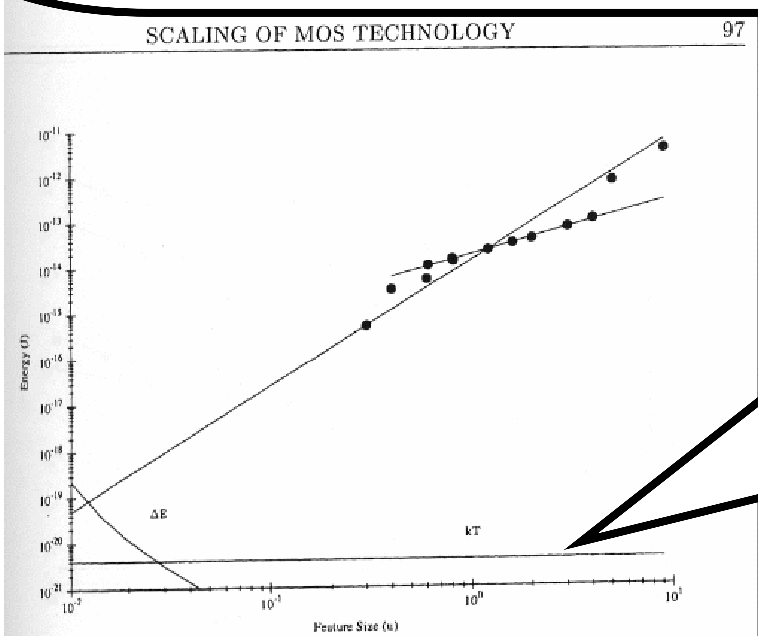

kT "helper line," drawn out of the reader's focus because it wasn't important at the time of writing
– Carver Mead, Scaling of MOS Technology, 1994

Sandia National Laboratories

Metaphor: FM Radio on Trip to Albuquerque

- You drive to Albuquerque 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

Sandia National Laboratories

FM Radio and End of Moore's Law

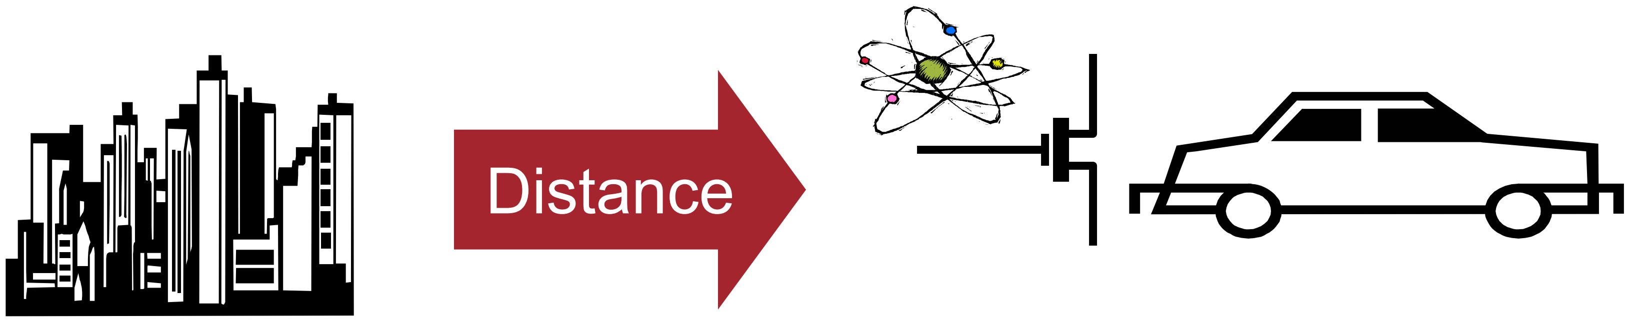

Driving away from FM transmitter→less signal
Noise from electrons → no change

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

Sandia National Laboratories

Personal Observational Evidence

- Have radios become better able to receive distant stations over the last few decades with a rate of improvement similar to Moore's Law?
- You judge from your experience, but the answer should be that they have not.
- Therefore, we infer that electrical noise does not scale with Moore's Law.

SIA Semiconductor Roadmap

- Generalization of Moore's Law
  - Projects many parameters
  - Years through 2016
  - Includes justification
  - Panel of experts
    - known to be wrong
  - Size between Albuquerque white and yellow pages

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, $I_{sd,leak}$ (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-μm) [9]	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} / $I_{dd\text{-}NMOS}$)(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

Scientific Supercomputer Limits

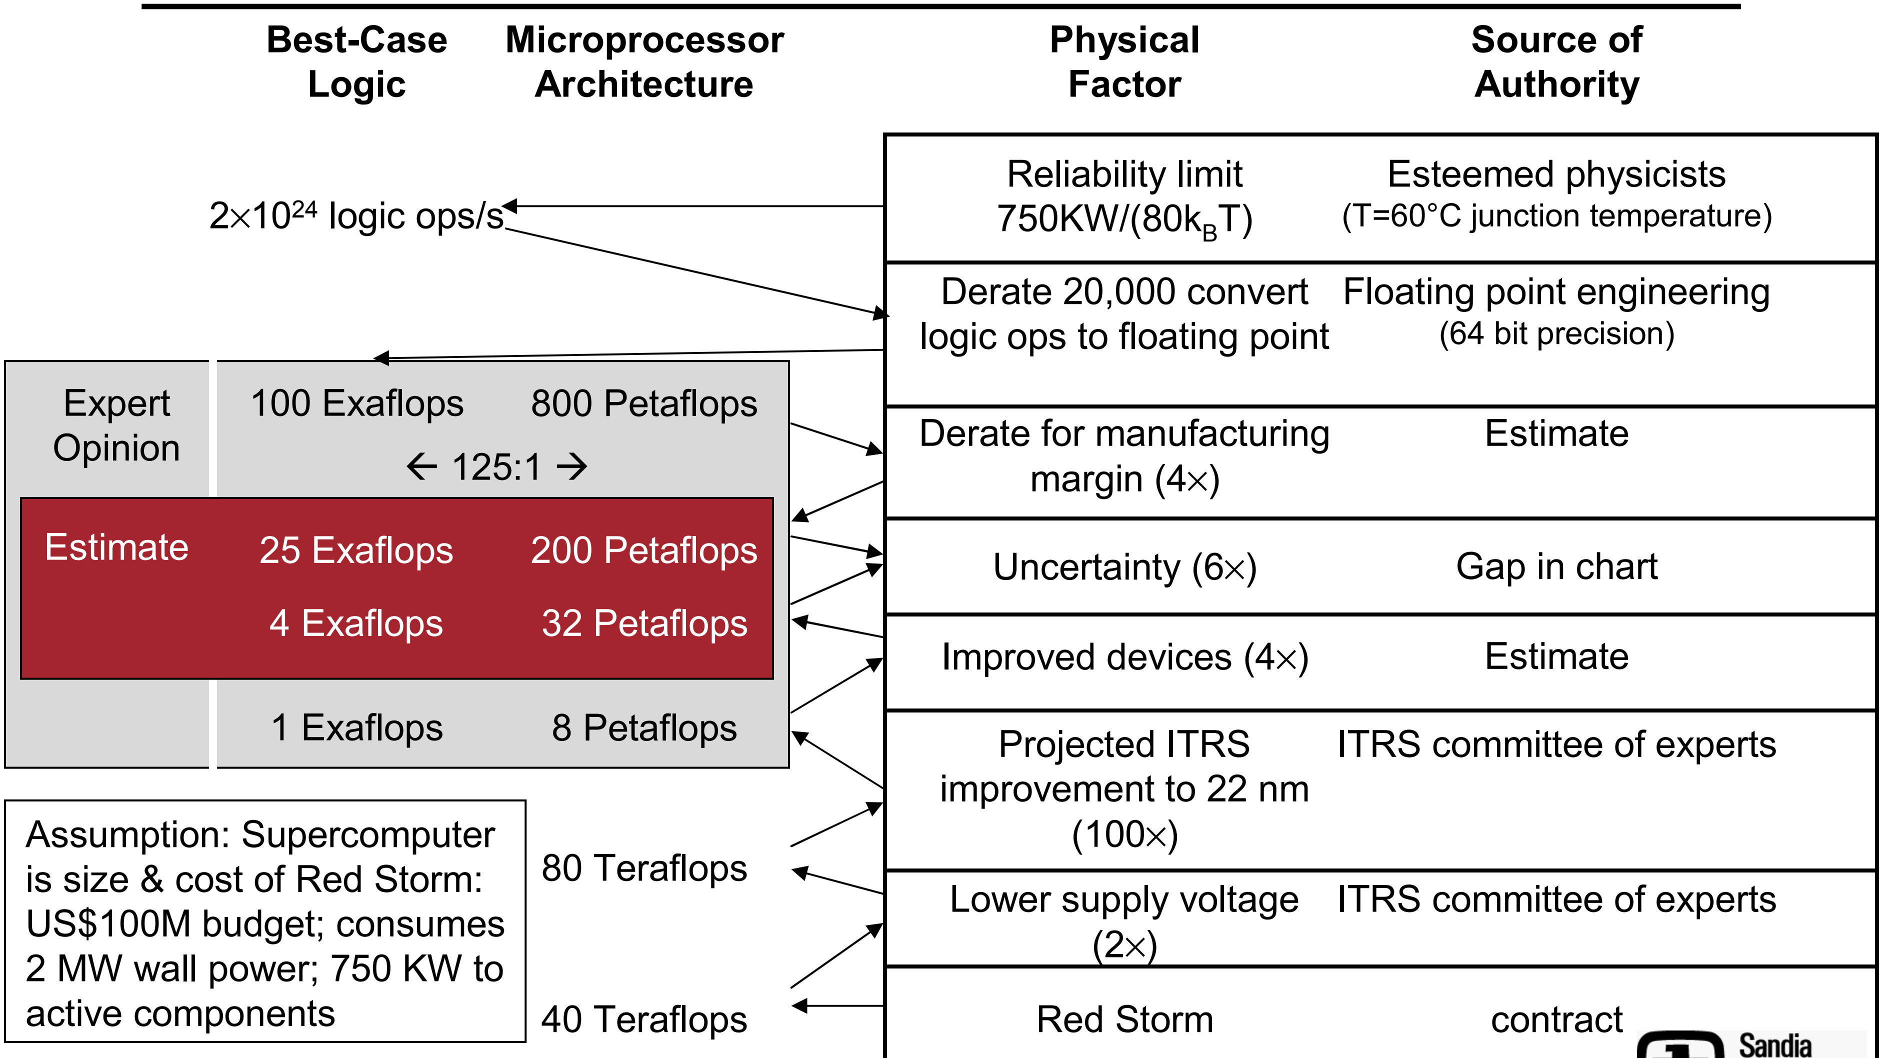

Sandia National Laboratories

Limits As Reported in ITRS Roadmap

- Track 1: Transistors
  - "Booster" innovations to improve transistors
    - Wrap-around gates
    - Different materials
    - … to total of 7 steps
  - Spreadsheet and MASTAR
    - Assume Moore's Law maintained
    - Solve for schedule
    - Uses 6/7 steps
- Track 2: Beyond Transistors
  - Broadbased study of "things that compute"
  - Reported power levels generally above $100k_BT$
  - Ignores "reversible logic" without explanation
  - Architectures not focused on supercomputing

Sandia National Laboratories

ITRS Transistor Geometries

- [Workshop participants reviewed pages 4 and 5 of the ITRS 2003 Emerging Devices Section. We do not have copyright permission to reproduce these pages.]
- Discussion was about changes in transistor geometry that lead to higher performance.

ITRS Technology Progression

- [Workshop participants reviewed page 12 of the ITRS 2003 Emerging Devices Section. We do not have copyright permission to reproduce this page.]
- Discussion was about how by 2016 6 of 7 performance boosters will have been used.

Sandia National Laboratories

ITRS Device Review 2016

- [Workshop participants reviewed page 42 of the ITRS 2003 Emerging Devices Section. We do not have copyright permission to reproduce this page.]
- Discussion was about how power performance of future proposed by ITRS was not all that much better than current transistors (with the exception of biologically-inspired and quantum devices, which were deemed very far out).

Conclusions on Limits

- Industry's roadmap is clear to $100k_BT$ limit
- Obvious step-up sequence for supercomputing
  - Follow Moore's Law to $100k_BT$
  - Advanced architectures good for 125×
  - I find no evidence that industry has collective direction beyond this
- Wait until this afternoon
  - Craig Lent's Quantum Dots include graphs showing 10^{-5} to 10^{-7} k_BT
  - You decide if he's right!

Sandia National Laboratories