Sandia National Laboratories report SAND2004-5872C
Unclassified Unlimited Release

Editor's note: These were presented by Erik DeBenedictis to organize the workshop

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.

Overall Motivation

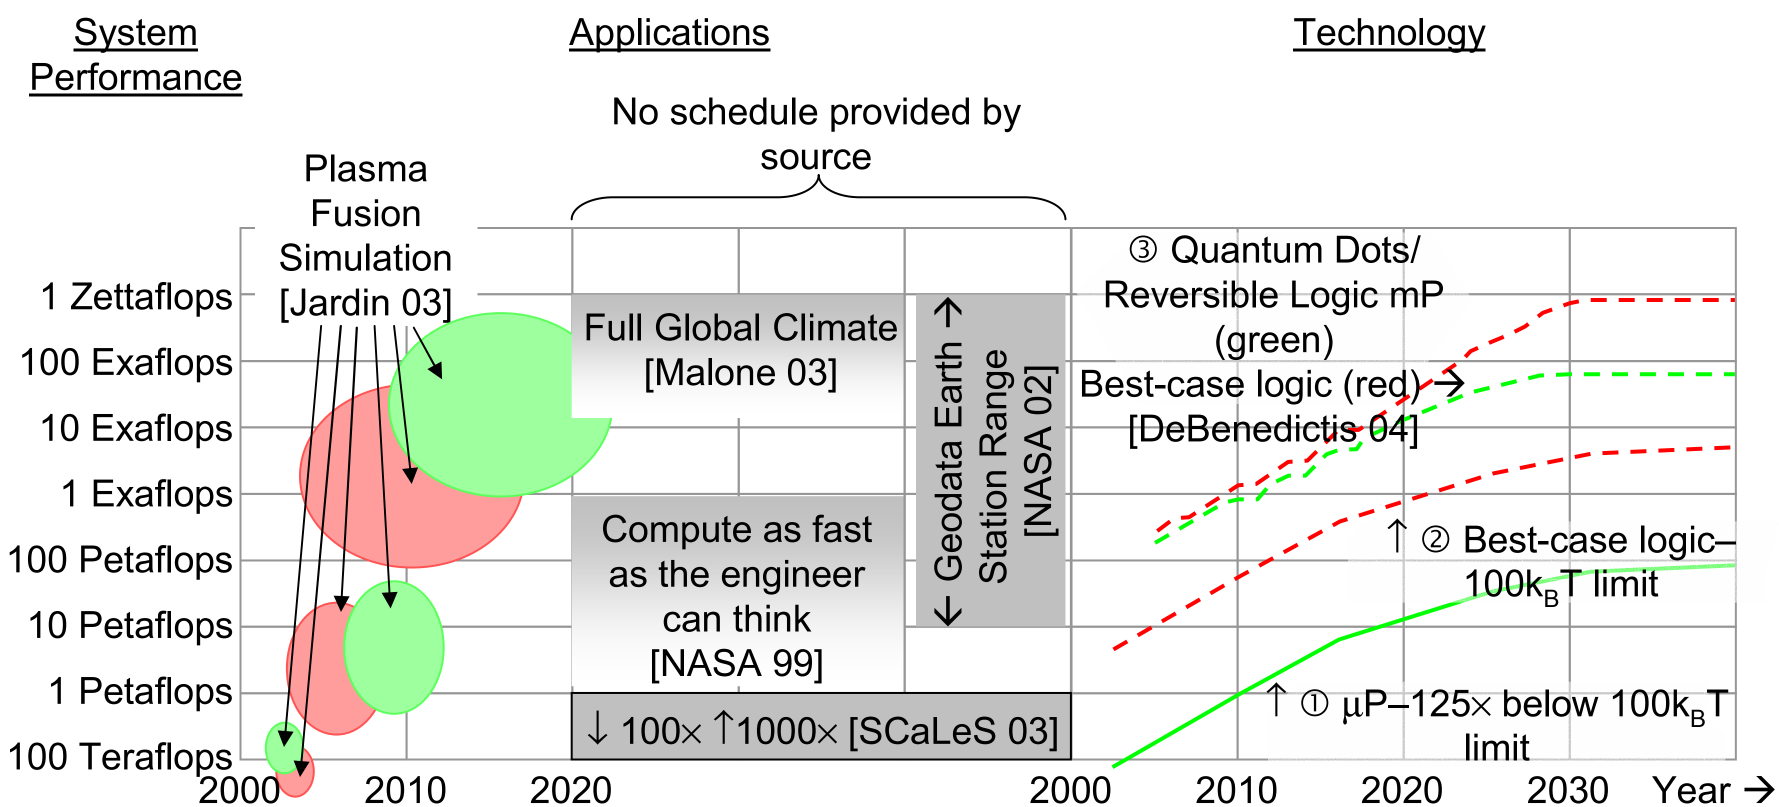

[Jardin 03] S.C. Jardin, "Plasma Science Contribution to the SCaLeS Report," Princeton Plasma Physics Laboratory, PPPL-3879 UC-70, available on Internet.
[Malone 03] Robert C. Malone, John B. Drake, Philip W. Jones, Douglas A. Rotman, "High-End Computing in Climate Modeling," contribution to SCaLeS report.
[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.
[NASA 02] NASA Goddard Space Flight Center, "Advanced Weather Prediction Technologies: NASA's Contribution to the Operational Agencies," 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.

The Back and Forth of Computing Limits

- 1. Public: Moore's Law continues forever
  - Justification: California real estate prices go up
  - Stock prices go up and up (in the late 1990s)
- 2. Industry (ITRS): Moore's Law ends
  - Justification: The types of technology my management is currently investing in is limited to level ____
- 3. 998 of 1000 physicists, including a dozen with Nobel prizes: No upper limit on computing per watt
  - Constructive solution: Reversible Logic
- 4. A couple skeptical physicists: Nobody has demonstrated devices below the k_BT limit
  - Could there be an undiscovered physical law?

Sandia National Laboratories

Morning Session

- Organizational
  - 9:00 Rob Leland
    - Host comments
  - 9:05 Erik DeBenedictis
    - Workshop organization
- A Big Application
  - 9:15 Philip Jones
    - Climate Modeling
- Current Technology Limits
  - 10:00 Erik DeBenedictis
    - ITRS Roadmap
- Break
- Advanced Architecture
  - 11:00 Peter Kogge
    - PIM architecture
- Software
  - 11:45 Bill Gropp
    - Software

Afternoon Session

- Logic
  - 2:00 Michael Frank
    - Reversible Logic
- Post Transistor Devices
  - 2:45 Craig Lent
    - Quantum Dots

- 4:00 Panel Session
  - Thomas Sterling, Caltech/JPL
  - Horst Simon, LBL/NERSC
  - David Koester, MITRE/DARPA HPCS
  - Terry Michalske, Center for Integrated NanoTechnology
  - Fred Johnson, DOE
  - Rob Leland, Sandia

Sandia National Laboratories

Climate Modeling

- About the Speaker
  - Philip Jones, Project Leader of the Climate Ocean and Sea Ice Modeling (COSIM), Los Alamos National Laboratory
  - Phil was part of the SCaLeS report study on computational requirements for climate ← modeling
    - See link on http://www.zettaflops.org

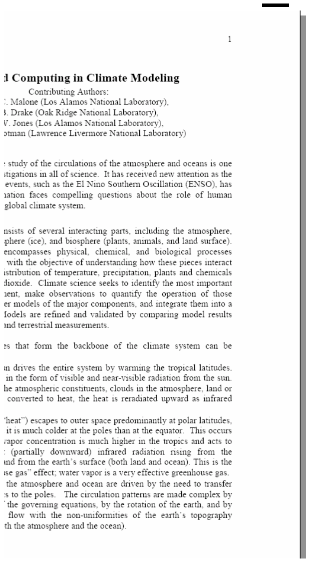

1

d Computing in Climate Modeling

Contributing Authors:
:. Malone (Los Alamos National Laboratory),
3. Drake (Oak Ridge National Laboratory),
V. Jones (Los Alamos National Laboratory),
otman (Lawrence Livermore National Laboratory)

: study of the circulations of the atmosphere and oceans is one stigations in all of science. It has received new attention as the events, such as the El Nino Southern Oscillation (ENSO), has iation faces compelling questions about the role of human global climate system.

nsists of several interacting parts, including the atmosphere, phere (ice), and biosphere (plants, animals, and land surface). encompasses physical, chemical, and biological processes with the objective of understanding how these pieces interact istribution of temperature, precipitation, plants and chemicals dioxide. Climate science seeks to identify the most important ient, make observations to quantify the operation of those er models of the major components, and integrate them into a Iodels are refined and validated by comparing model results and terrestrial measurements.

es that form the backbone of the climate system can be

in drives the entire system by warming the tropical latitudes. in the form of visible and near-visible radiation from the sun. he atmospheric constituents, clouds in the atmosphere, land or converted to heat, the heat is reradiated upward as infrared

'heat") escapes to outer space predominantly at polar latitudes, it is much colder at the poles than at the equator. This occurs vapor concentration is much higher in the tropics and acts to : (partially downward) infrared radiation rising from the ind from the earth's surface (both land and ocean). This is the ise gas" effect; water vapor is a very effective greenhouse gas. the atmosphere and ocean are driven by the need to transfer :s to the poles. The circulation patterns are made complex by the governing equations, by the rotation of the earth, and by flow with the non-uniformities of the earth's topography th the atmosphere and the ocean).

- Notes
  - A very important problem for humanity
  - Many levels of increasing sophistication to 1 ZFLOPS.
  - Independent validation
    - NASA study for ground processing of Earth sciences data
    - Challenge questions

NASA Climate Earth Station

Based on these inputs, various portions of the Modeling and Data Assimilation System will require anywhere from 10^7 to 10^{13} GFLOPS of computational resources. In other words, the range of computational resources needed is 10^{16} to 10^{21} Floating Point Operations per Second. For the curious, the range can also be stated as 10 PetaFLOPS to 1 ZettaFLOPS.

4.1.2. Anticipated Computing Technology Capabilities

At first glance, the numbers discussed in the previous section appear so high as to be impossibly ludicrous. However, with the expected growth in computing capabilities, the lower end of this spectrum actually falls within the domain of possibility.

– **"Advanced Weather Prediction Technologies: NASA's Contribution to the Operational Agencies," Gap Analysis Appendix, May 31, 2002**

Sandia National Laboratories

Technology Limits

- About the Speaker
  - Erik DeBenedictis, staff member at Sandia National Laboratories
- Notes
  - The "limits of current technology" involves two questions
    - The social question of the dividing line between "current" and "future" technology
    - The technical limits of the technology

Sandia National Laboratories

Advanced Architectures

- About the Speaker
  - Peter Kogge, Professor at Notre Dame, first holder of the endowed McCourtney Chair in Computer Science and Engineering (CSE)
  - IBM Federal Systems Division, from 1968 until 1994
  - IEEE, IBM fellow
  - Ph. D. Stanford, 1973
- Notes
  - Advanced architectures such as PIM appear to be the first option beyond simple continuation of Moore's Law
  - Upside potential about 100×

Sandia National Laboratories

Software

- About the speaker
  - Bill Gropp is Associate Division Director, Senior Computer Scientist, Mathematics and Computer, Science Division, Argonne National Laboratories
  - Ph. D. Stanford 1982
  - Well known for MPI
- Notes:

Reversible Logic

- About the Speaker
  - Michael Frank, Assistant Professor, Florida State University
  - Ph. D. MIT 1999, “Reversibility for Efficient Computing”
- Notes
  - Reversible logic is essential to beat the limits Erik described
  - It works by recycling energy instead of turning it into heat
  - Reversible logic is widely accepted in the physics community, but not broadly understood

Quantum Dots

- About the Speaker
  - Craig Lent, Freimann Professor of Engineering, University of Notre Dame
- Notes
  - Quantum Dots for computation are a promising device technology that could reach to Zettaflops
  - Published material on quantum dots is mature enough to estimate logic performance, and this performance is pretty good

Panel Session and Group Process

- Panel Session
  - Question: “How much should we change supercomputing to enable the applications that are important to us, and how fast.”
- Results of Workshop
  - Each panelist may propose one or two concluding statements
    - “Moore’s Law will/won’t solve all problems if you wait long enough”
  - Audience will vote
  - Statements and degree of agreement will be the conclusion of workshop

Organizer’s View of Appropriate Answers*

*Editor’s note: Which were generally ignored

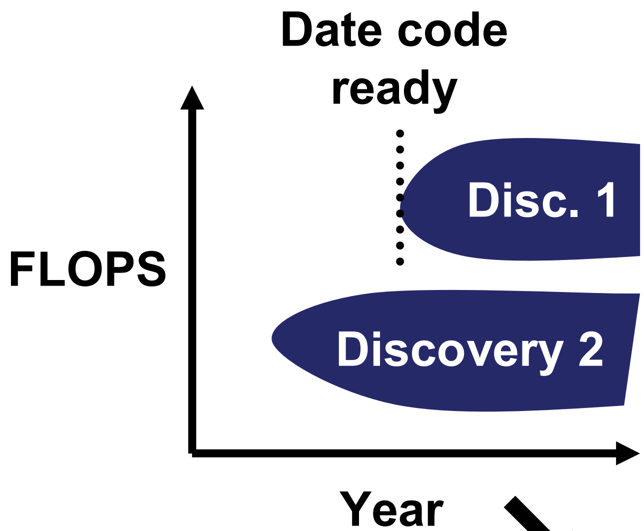

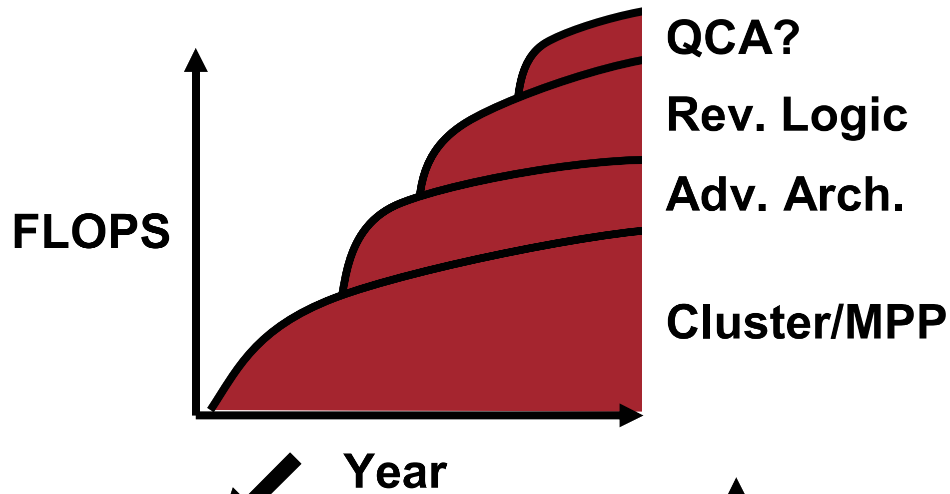

Notes:
* Not necessarily one machine; different applications may require different machines
* Specifics are just my ideas

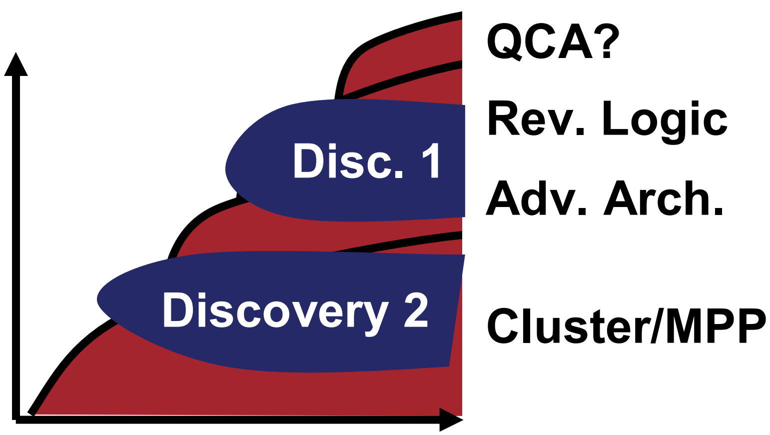

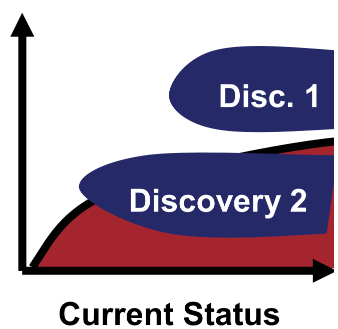

Sandia National Laboratories

Thomas Sterling

- Zettaflops at nano-scale technology is possible [Vote: 11/22]
  - Size requirements tolerable
    - But packaging is a challenge; [Vote: 21/21]
- Major obstacles [Vote for only 2 of the options below]
  - Power [Vote: 13/22]
  - Latency [Vote: 0/22]
  - Parallelism [Vote: 12/22]
  - Reliability [Vote: 4/22]
  - Programming [Vote: 8/22]

Horst Simon

- A Zettaflops computer will have emergent intelligent behavior. [Vote: 8/22]
- The first sustained Petaflops application that wins the G. Bell award will use MPI. [Vote: 12/22]
- The first sustained Exaflops application that wins the G. Bell award will use MPI. [Vote: 2/22]

- [Editors note: The qualification of winning the Gordon Bell award implies a general purpose computer and software written in a high level language.]

Sandia National Laboratories

David Koester

- Moore's Law doesn't matter as long as we need to invest the increase in transistors into machine state — i.e., overhead — instead of real use
- Keep putting more transistors out there and power stays the same (130 W/chip) all going into overhead; work is only increasing a little by clock rate.
- Moore's law doesn't translate into increase in value.
- [Vote for all above: 6 agree, 2 disagree]

Terry Michalske

- Will we always have a fixed architecture that we put an operating system onto…or will the architecture (hardware) reconfigure itself to run the application [Editors note: software identifies parts of hardware that have faults and configures itself to avoid these parts. Vote: 15/22]

Sandia National Laboratories

Fred Johnson

- Need to also consider multi-scale math, as a new initiative. [Vote: 21/21]
- Yet again I/O has been left in the closet [Vote: 21/21]

Rob Leland

- **Early speakers take all the time presenting, so the last speaker need not have slides.** [Editor's note: …and so Rob didn't ☺. Vote: 21/21]
- **Do we want to be leaders in changing supercomputing?** [Vote: 21/21]
- **Do we want to create possibilities that have not been imagined?** [Vote: 21/21]
- **How aggressively do we go after Zettaflops? It takes 20 years to develop a supercomputing technology to full production. If the case for Zettaflops in 2025 is strong, we need to start now. Statement: "It is exactly the right time to start thinking about Zettaflops seriously."** [Vote: 21/21]

Sandia National Laboratories

Erik DeBenedictis

- The workshop articulated a series of constructive steps toward very high performance computers, at the level of Zettaflops. [Vote: 10/21]
- The issues deserve more investigation [Vote: 21/21]
- Reversible logic needs a more thorough understanding [Vote: 15/21]

Sandia National Laboratories