Enabling technologies for beyond exascale computing

Paul Messina

Director of Science

Argonne Leadership Computing Facility

Argonne National Laboratory

July 9, 2014

Cetraro

Do technologies cause revolutions in computing, as opposed to evolution?

- **The transistor**
  - Enabled much bigger systems -- not practical to have billions of vacuum tubes
  - The transistor was not invented with the goal of finding a replacement for vacuum tubes
  - Is used both for logic and for memory
- **But in what sense are computers built using transistors a revolution over those that use vacuum tubes?**
- **At high level, architectures did not change**

Different technologies: analog vs. digital

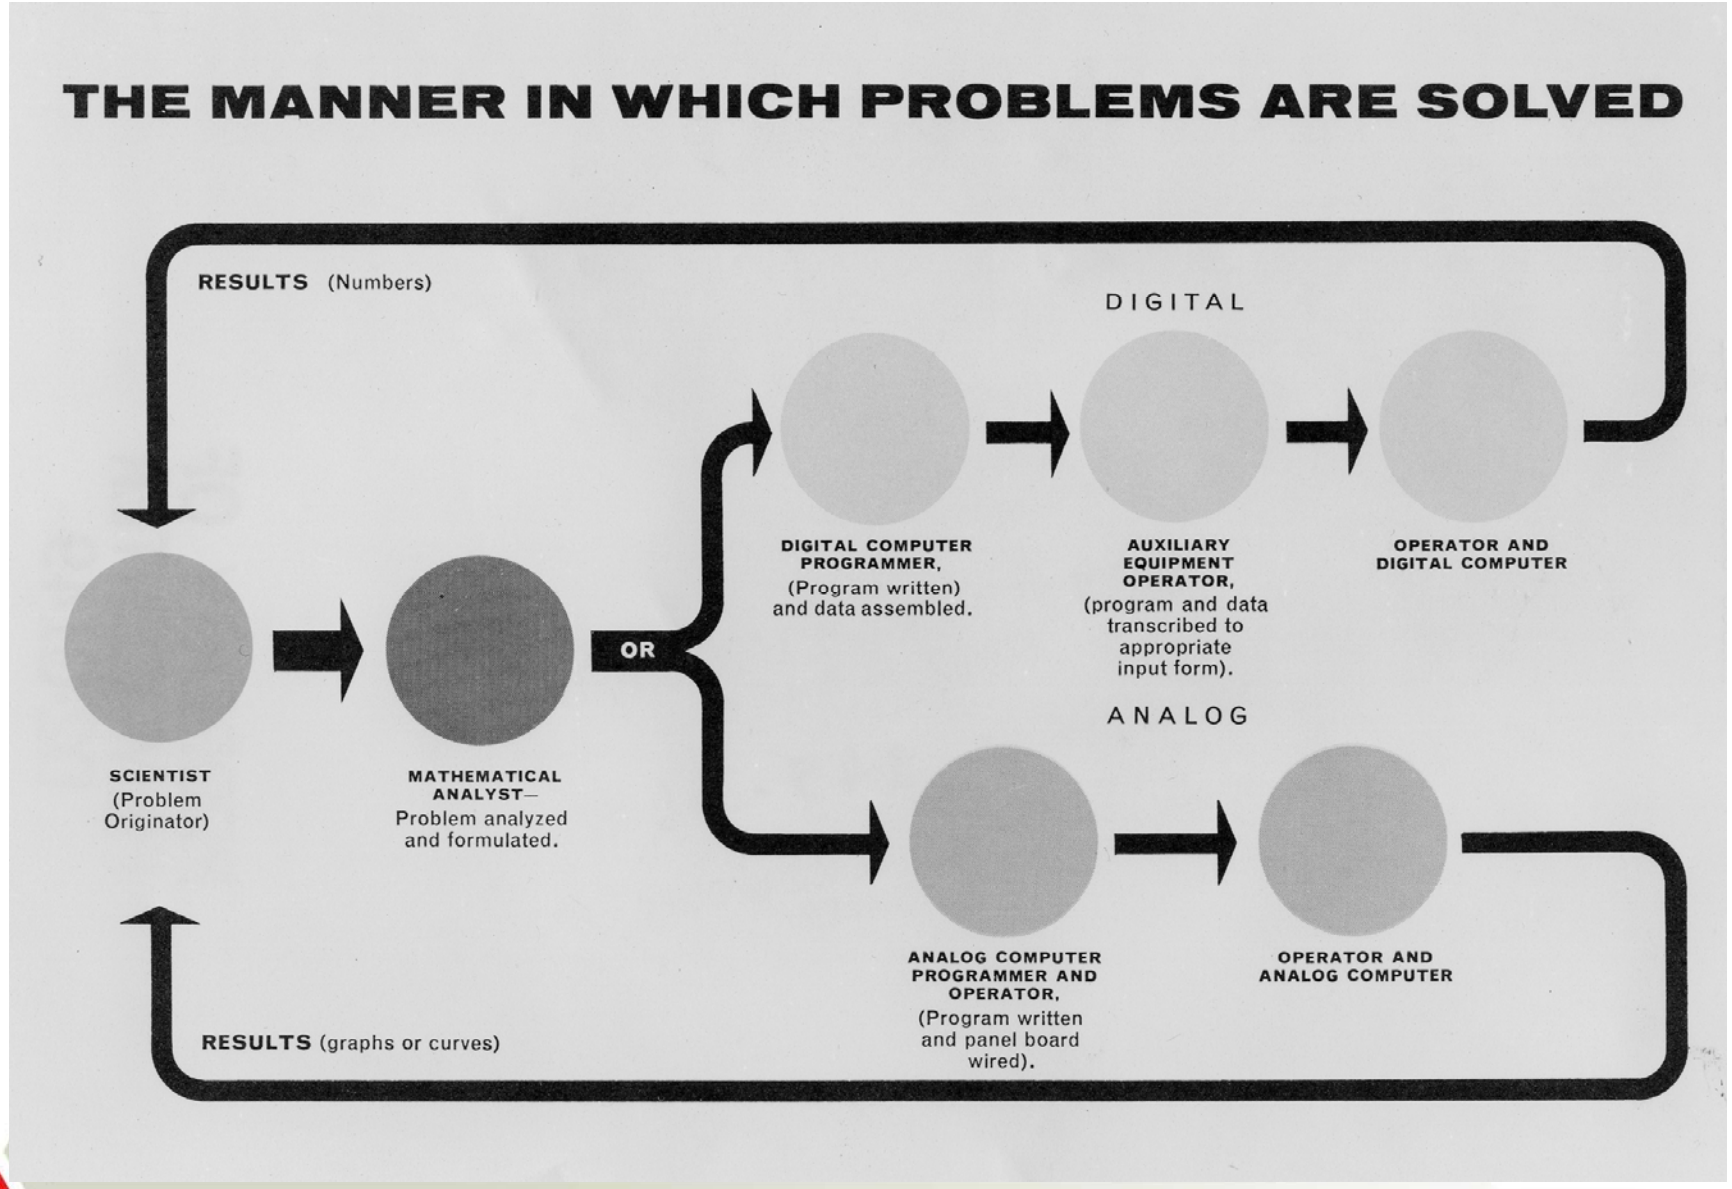

PACE Analog Computer, 1964
Argonne National Laboratory, Frank Morehouse

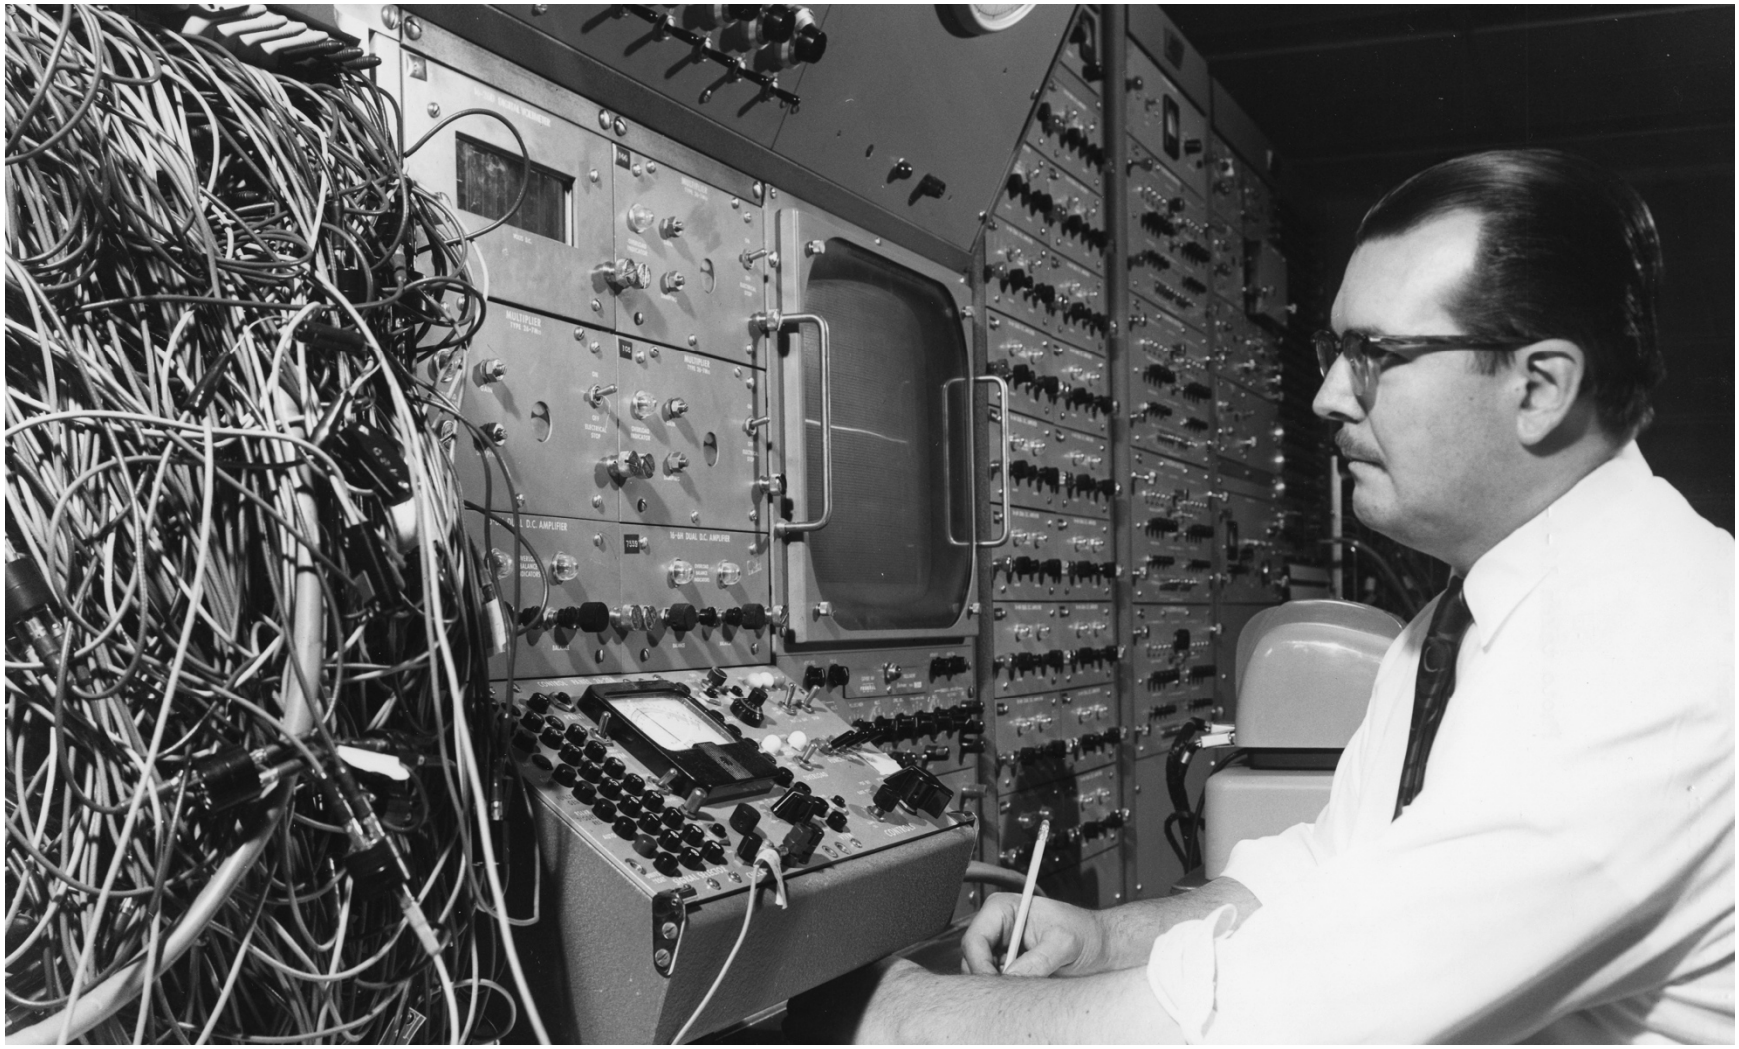

An analysis of Analog vs. Digital, c. 1968

A COMPARISON OF THE USE OF ANALOG AND DIGITAL COMPUTERS

FOR THE SOLUTION OF DIFFERENTIAL EQUATIONS*

Nye F. Morehouse, Jr. and Alan Winiecki
Argonne National Laboratory, Argonne, Illinois

SUMMARY

A comparison is made between the use of analog and digital computers for the solution of systems of ordinary differential equations. The use of continuous difference equations for estimating the truncation error is illustrated and some general criteria developed for the number of mesh points per cycle required in digital methods. Although only linear systems are considered, the method can be extended to non-linear systems.

Distributed memory and address space

Caltech Cosmic Cube

A Beowulf cluster

CASA Gigabit Testbed 1989-1993

- One of first 5 gigabit network testbeds established by DARPA/NSF in 1989
- First to focus on meta-supercomputer applications in seismic image processing, global climate modeling, and chemical reaction dynamics
- Applications and network engineering team: Caltech, JPL, SDSC, LANL
- SONET OC-48 transport provided by: MCI, PacBell, and US West

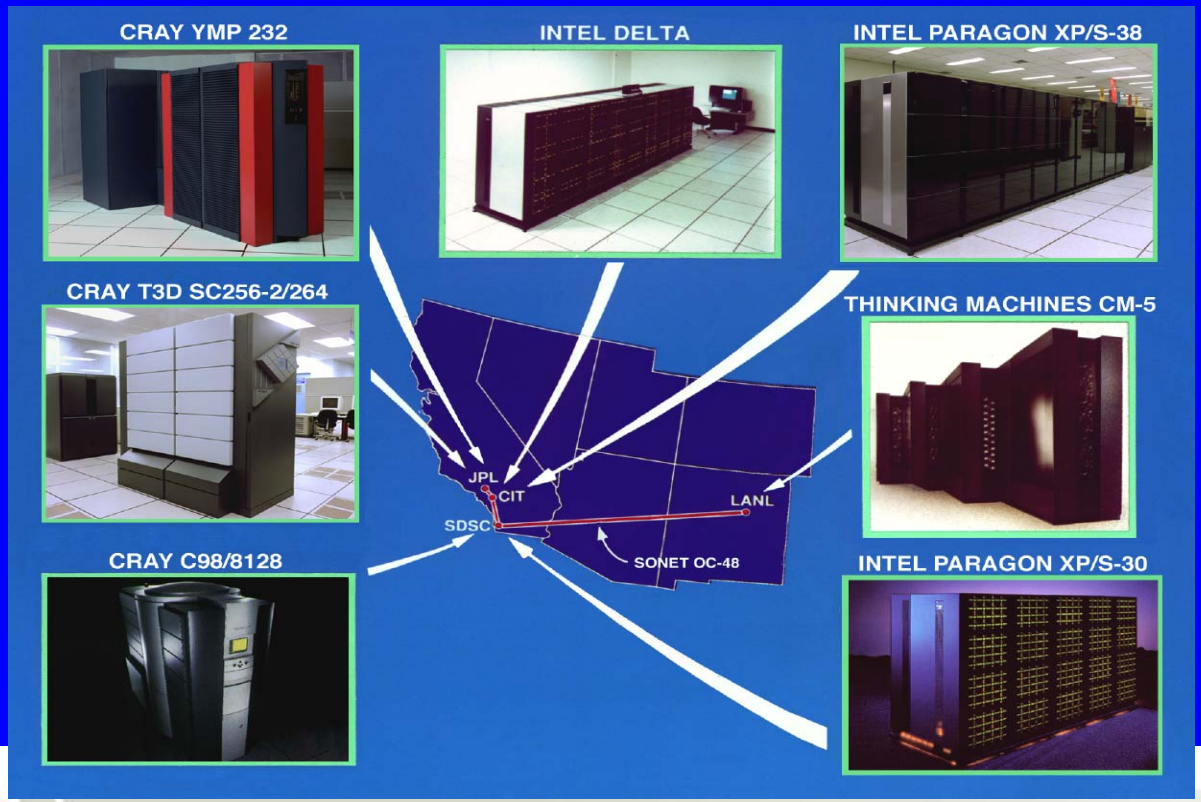

Salient Features

- First to demonstrate non-linear algorithm speedup with meta-supercomputer
- 800 Mbit/s
- 300 GFLOPS
- 2000 km long

3D View of Turtle Mountains, CA
Crustal Block Diagram of Geophysical Data

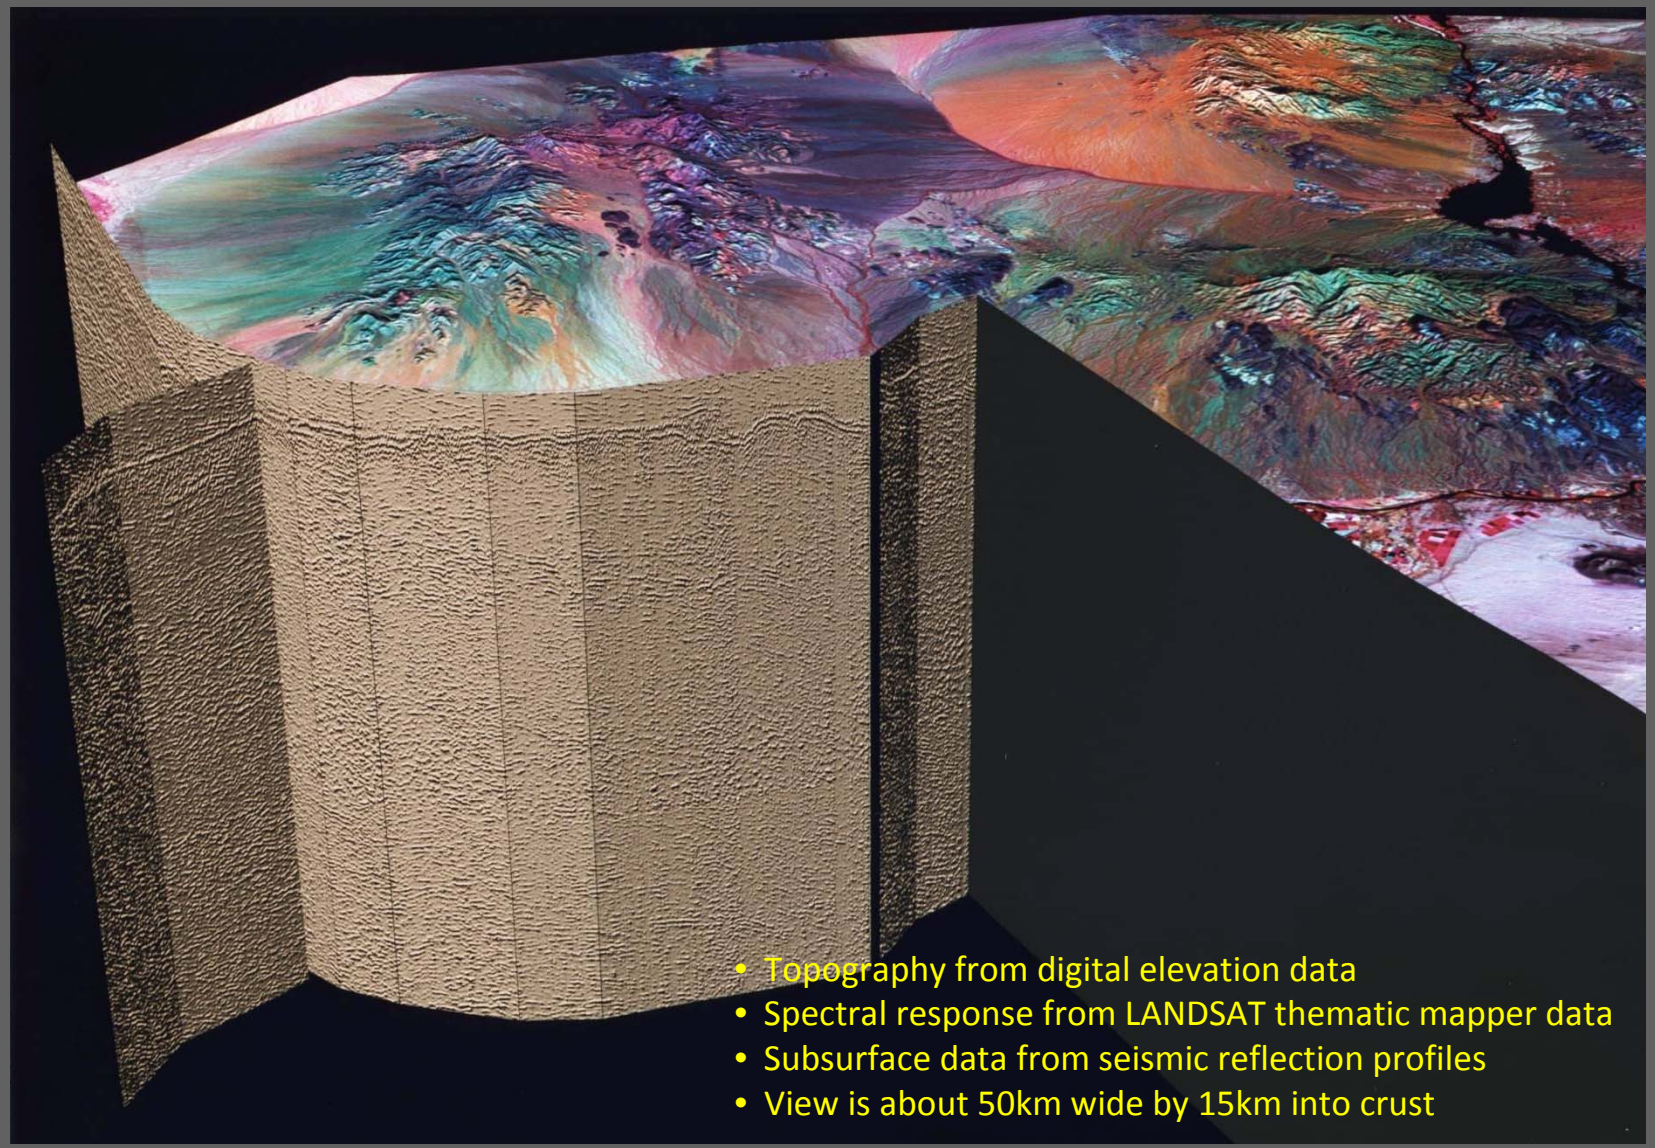

Foster and Kesselman abstracted the CASA metacomputing approach, renamed it “the grid,” and produced a more refined way to do that type of computing: GLOBUS

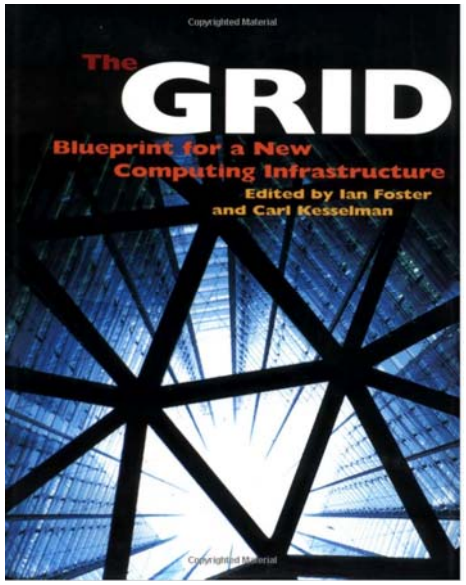

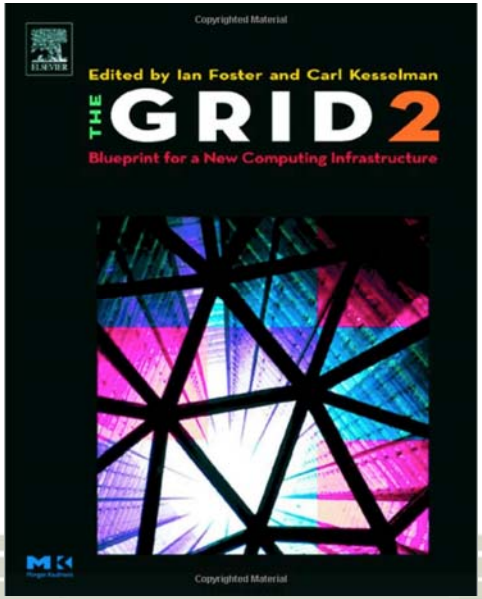

“The grid promises to fundamentally change the way we think about and use computing. This infrastructure will connect multiple regional and national computational grids, creating a universal source of pervasive and dependable computing power that supports dramatically new classes of applications.”

“The Grid is an emerging infrastructure that will fundamentally change the way we think about - and use - computing.”

And now we have

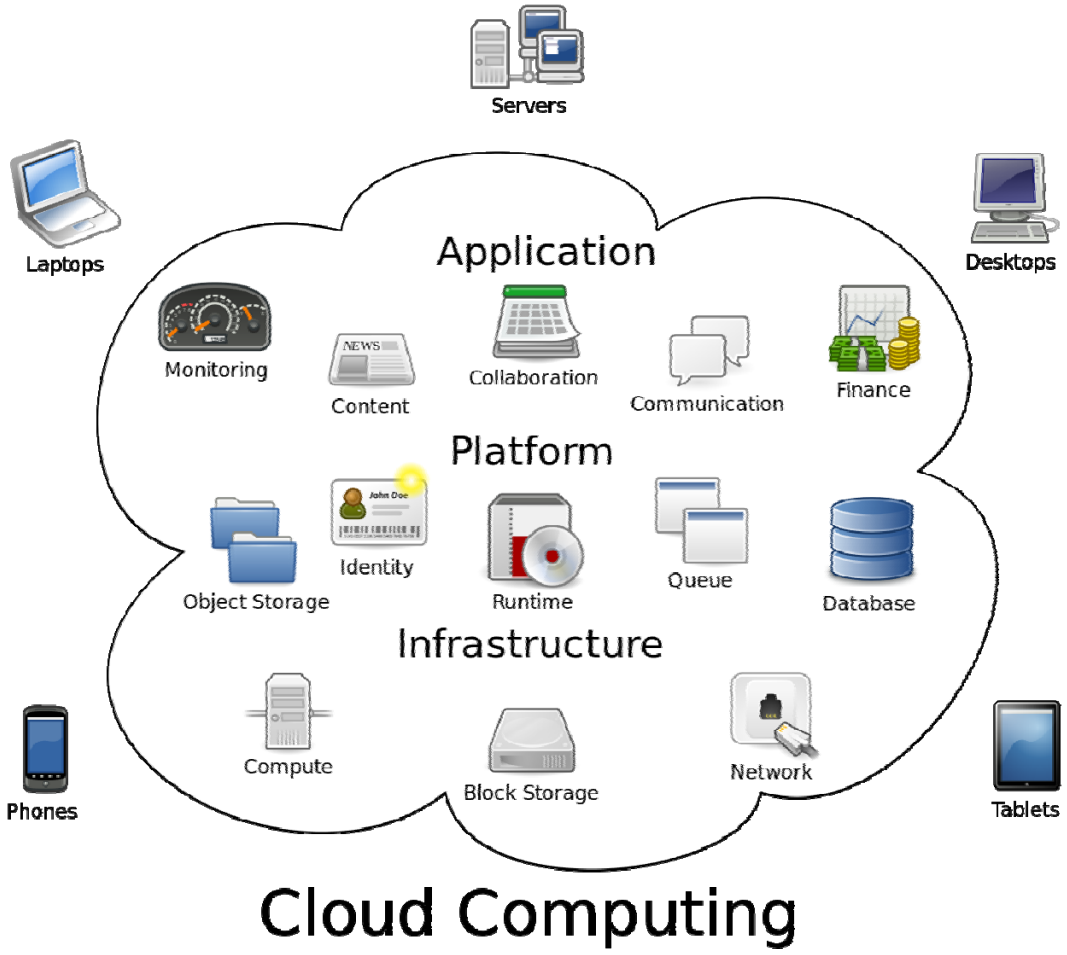

Highly distributed address space

First Workshop on "Enabling Technologies for Petaflops Computing"
February 22-24, 1994 - Pasadena, California

A past attempt to predict and plan the future

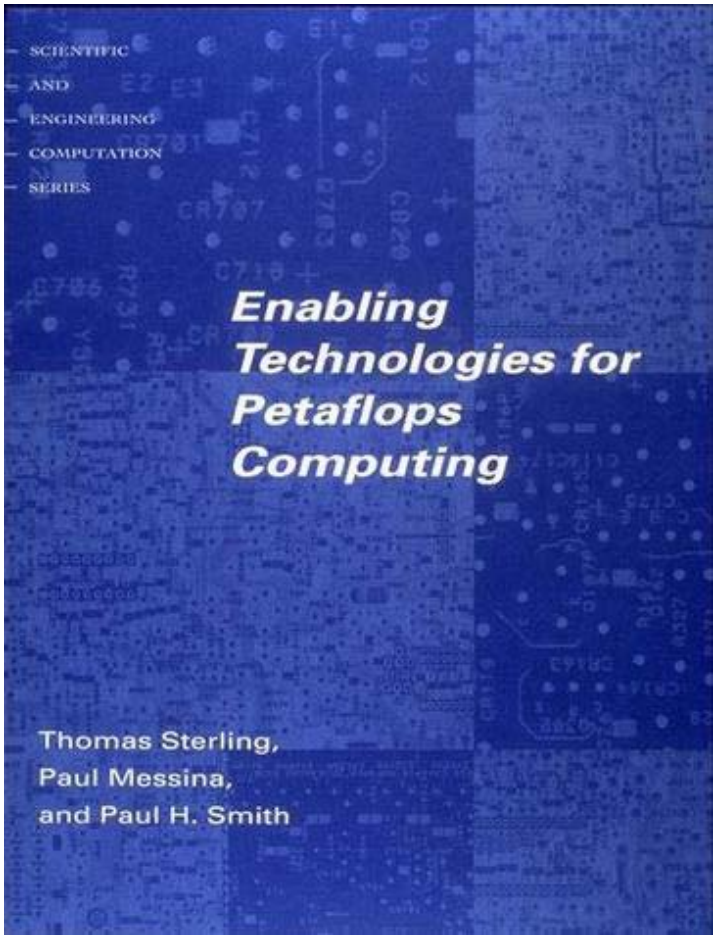

- The workshop focused on four interrelated areas: applications and algorithms, device technology, architecture and systems, and software technology.
- "Construction of an effective petaFLOPS computing system will be feasible in two decades, although effectiveness and applicability will depend on dramatic cost reductions as well as innovative approaches to system software and programming methodologies; a mix of technologies such as semiconductors, optics, and possibly cryogenics will be required"

Alas, it is always dangerous to prophesy, particularly, as the Danish proverb says, about the future.

- **But the workshops on petaflops did yield useful insights and a number of predictions that turned out to be on target, such as**
- **Memory size scaling much less than FLOPS will be adequate for applications that have N**3/4 memory size scaling**
- **Hybrid node architectures with different “size” units**
- **0.5 MW per PF in 2014**

What is “Beyond Exascale Computing?”

- We are not referring to 10**21 flops
- “Beyond exascale” systems as we are defining them will be based on new technologies that will finally result in the much anticipated (but unknown) phase change to truly new paradigms/methodologies. The session will therefore also include presentations on architecture advances that may be enabled as a consequence of technology progress.
- The focus of this session is principally on forward-looking technologies that might determine future operational opportunities and challenges for computer systems beyond the exascale regime.

Today's talks

- Rick Stevens, Argonne and University of Chicago: *"Beyond Exascale — What will Sustain our Quest for Performance in a Post-Moore World?"*
- Mikhail Dorojevets, Stony Brook University: "*Energy-Efficient Superconductor Circuits for High-Performance Computing*"
- Matthias Troyer, ETH-Zurich: *"Quantum Computing"*
- Mark Moraes, D.E. Shaw Research: "*Scaling lessons from the software challenges in Anton, a special-purpose machine for molecular dynamics simulation*"
- Patrick DeMichel, Hewlett-Packard: "*New technologies that disrupt our complete ecosystem and their limits in the race to Zettascale*"
- Keren Bergman, Columbia U.: "*Scalable Computing Systems with Optically Enabled Data Movement*"
- Robert Wisniewski, Intel Corp.: "*System Software for PEZ(Y)*"
- Thomas Sterling, Indiana U.: "*Extreme-scale Architecture in the Neo-Digital Age*"
- Panel on Beyond Exascale Computing

Acknowledgments

- Lucio Grandinetti conceived this session and challenged me to implement his vision

- Fortunately for me I was able to convince Thomas Sterling and Bob Lucas to do most of the work

- I want to thank in advance the speakers and panelists for having the courage to participate in this session

Breakthroughs can be costly

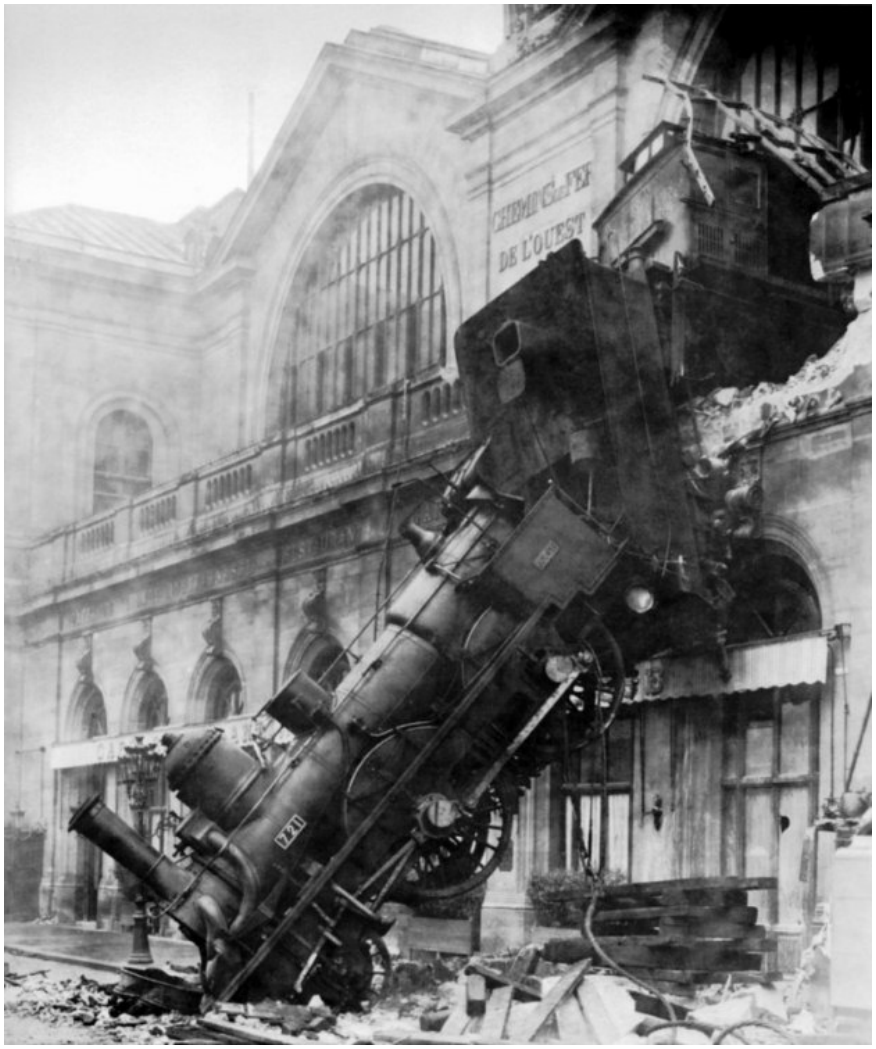

Goals for today's session

- **Stimulate investigation of potential sea change in computing based on**
  - characteristics of new and evolving technologies
  - Insights from the last 70 years, and
  - Speculations that do not violate too many laws of physics
- **In other words, the beginnings of an attempt to predict *and plan* the future**

Thank you

January 15, 1953, Washington DC