&

CRAY

Through the years…

When did it all begin?

1974?

1978?

1963?

CDC 6600 – 1974 NERSC started service with *the* first Supercomputer…

- A well-used system - Serial Number 1
  - On its last legs…
- Designed and built in Chippewa Falls
- Launch Date: 1963
- Load / Store Architecture
  - First RISC Computer!
- First CRT Monitor
- Freon Cooled
- State-of-the-Art Remote Access at NERSC
  - Via 4 acoustic modems, manually answered capable of 10 characters /sec

50th Anniversary of the IBM / Cray Rivalry…

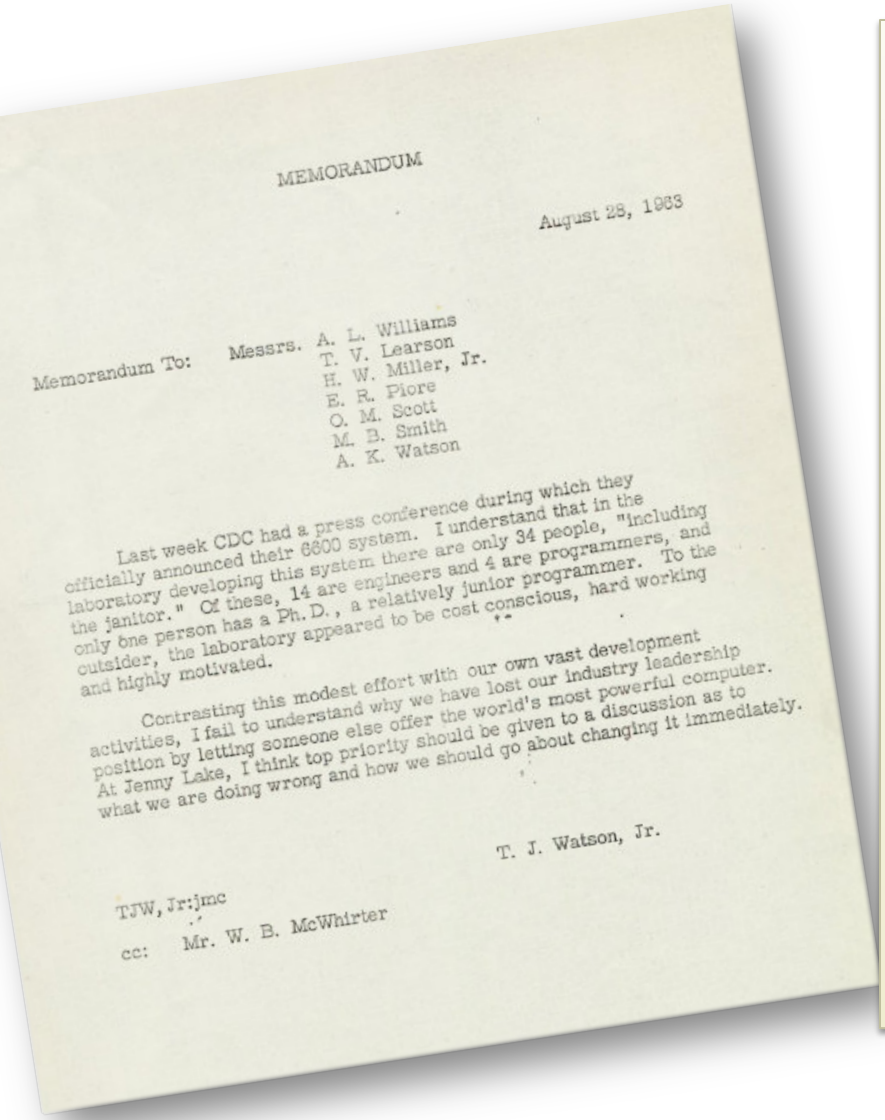

MEMORANDUM

August 28, 1963

Memorandum To: Messrs. A. L. Williams
T. V. Learson
H. W. Miller, Jr.
E. R. Piore
O. M. Scott
M. B. Smith
A. K. Watson

Last week CDC had a press conference during which they officially announced their 6600 system. I understand that in the laboratory developing this system there are only 34 people, "including the janitor." Of these, 14 are engineers and 4 are programmers, and only one person has a Ph.D., a relatively junior programmer. To the outsider, the laboratory appeared to be cost conscious, hard working and highly motivated.

Contrasting this modest effort with our own vast development activities, I fail to understand why we have lost our industry leadership position by letting someone else offer the world's most powerful computer. At Jenny Lake, I think top priority should be given to a discussion as to what we are doing wrong and how we should go about changing it immediately.

T. J. Watson, Jr.

TJW, Jr:jmc

cc: Mr. W. B. McWhirter

Last week, CDC had a press conference during which they officially announced their 6600 system. I understand that in the laboratory developing this system there are only 32 people, "including the janitor"…

Contrasting this modest effort with our vast development activities, I fail to understand why we have lost our industry leadership position by letting someone else offer the world's most powerful computer…

T.J. Watson, August 28, 1963

CDC 7600 – 1975

- Delivered in September
- 36 Mflop Peak
- ~10 Mflop Sustained
- 10X sustained performance vs. the CDC 6600
- Fast memory + slower core memory
- Freon cooled (again)

Major Innovations
- 65KW Memory
- 36.4 MHz clock
- Pipelined functional units

Cray-1 – 1978

- Serial 6
- An fairly easy transition for application writers
- LTSS was converted to run on the Cray-1 and became known as CTSS (Cray Time Sharing System)
- Freon Cooled (again)
- 2nd Cray 1 added in 1981

NERSC transitions users to vector architectures

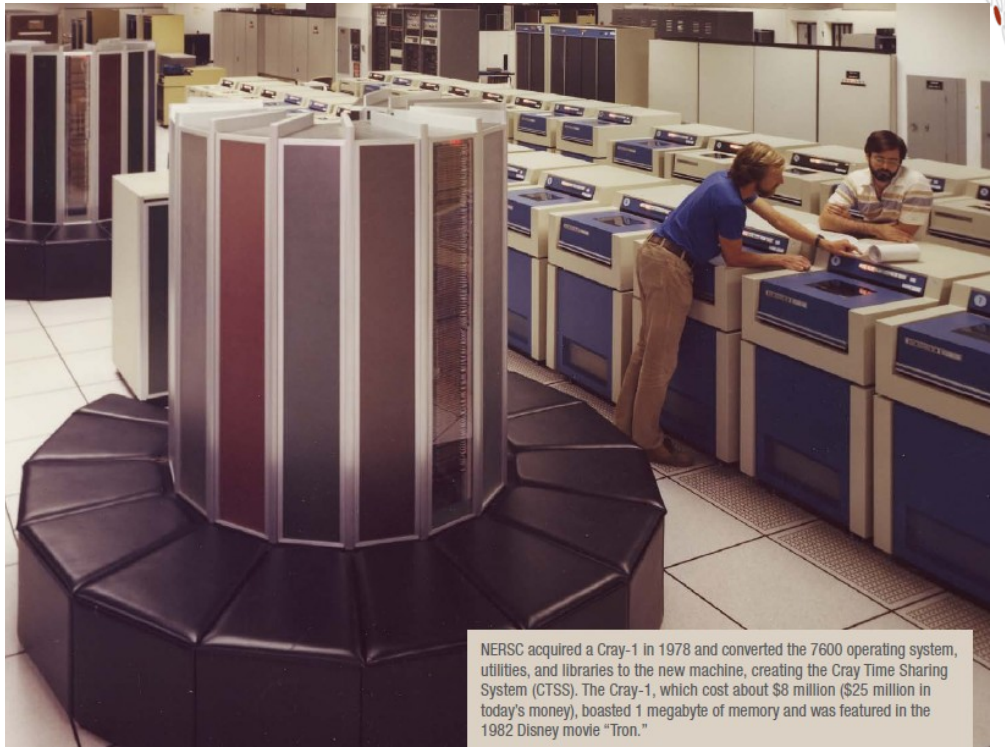

NERSC acquired a Cray-1 in 1978 and converted the 7600 operating system, utilities, and libraries to the new machine, creating the Cray Time Sharing System (CTSS). The Cray-1, which cost about $8 million ($25 million in today's money), boasted 1 megabyte of memory and was featured in the 1982 Disney movie "Tron."

Major Innovations
- Vector Processing
- Dependency Analysis in Compilers
- Integrated Circuits
- Packaging

Cray XMP added in 1984

- First multi-processor system
- System had some great architectural improvements
- At this point NERSC is working with 3500 users on their systems

NERSC transitions users to Multiple Processors ?

Major Innovations
- Multiple processors
- Multi-port memory
- Gather / Scatter
- Chaining
- CFT2

Cray-2- 1985

- Serial 1
- 4.1 ns clock
- 8-Processor Cray-2 delivered in 1990
- Best looking machine every built…

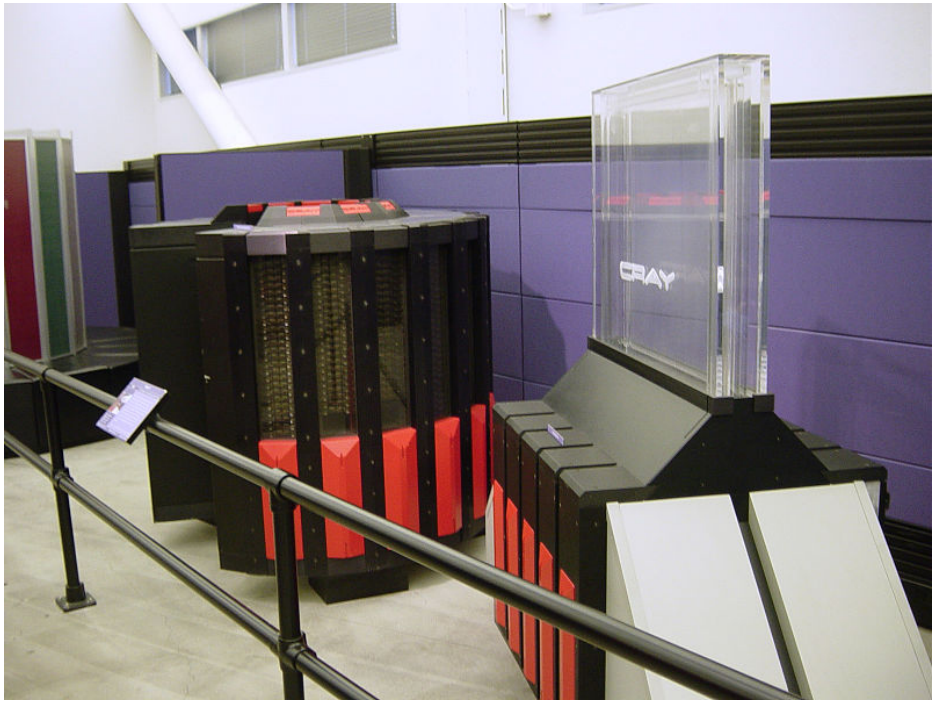

Major Innovations
- Packaging & Cooling
- Large memory
- Local Memory
- Macrotasking
- Microtasking

Cray C90 - 1992

- Interim Y-MP was installed until this could be delivered
- First non-classified installation of a C90
- Last “flagship” vector system at NERSC
- Top500 started in 1993 and this system was in position 19
  - It could have been #6


CRAY Y-MP C90 Computer System

Major Innovations
- 16 Processors
- Dual-pipe vectors
- Autotasking (OpenMP)

Back Then, We Fabbed our Own Chips…

Chippewa Falls Manufacturing Campus

Meanwhile, in the Marketplace, a New Style of Computing was Emerging…

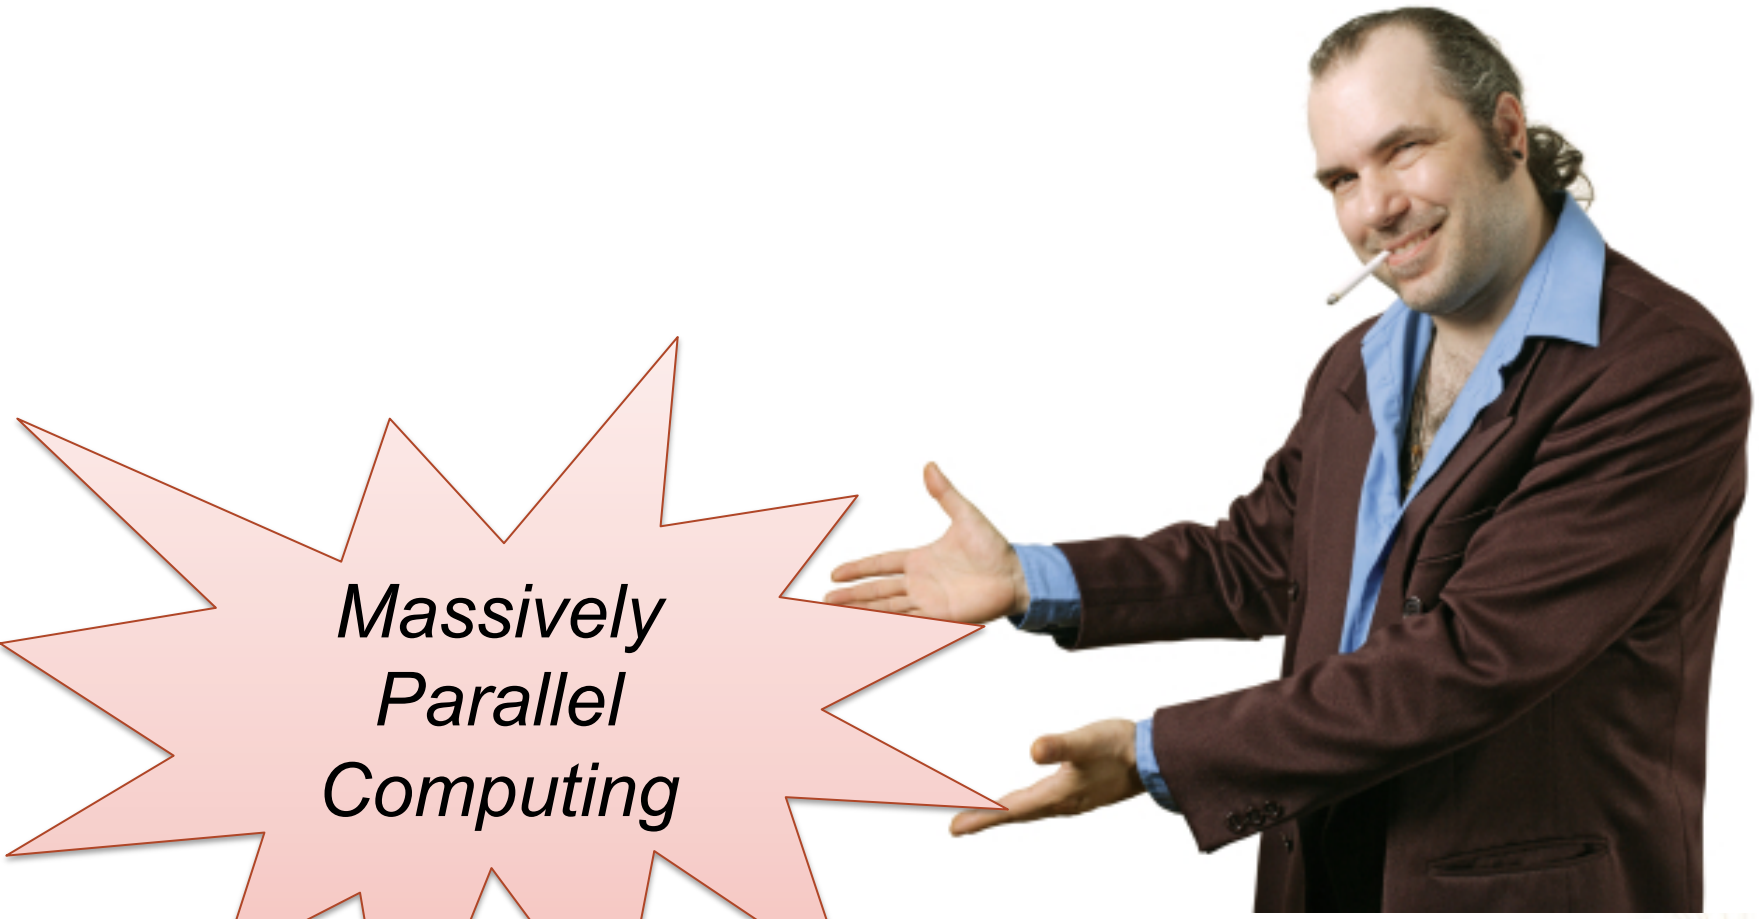

Can MPP Systems Really Perform ?

Twelve Ways to Fool the Masses When Giving Performance Results on Parallel Computers

David H. Bailey

June 11, 1991

Ref: Supercomputing Review, Aug. 1991, pg. 54--55

Abstract

Many of us in the field of highly parallel scientific computing recognize that it is often quite difficult to match the run time performance of the best conventional supercomputers. This humorous article outlines twelve ways commonly used in scientific papers and presentations to artificially boost performance rates and to present these results in the "best possible light" compared to other systems.

Cray T3D MPP Testbed - 1994

- Grew out of Steve Nelson's MPP Advisory Group
- 18 Months from concept to delivery
  - having your own fab can come in handy
- PATP – Parallel Application Technology Program
- 150 Mhz Alpha EV4
- Y-MP "front end" (the first esLogin ?)

Major Innovations
- One-sided get/puts
- Low Latency…
- 3D Torus
- Density (cold plate cooling)
- CRAFT (sort of)
- PVM

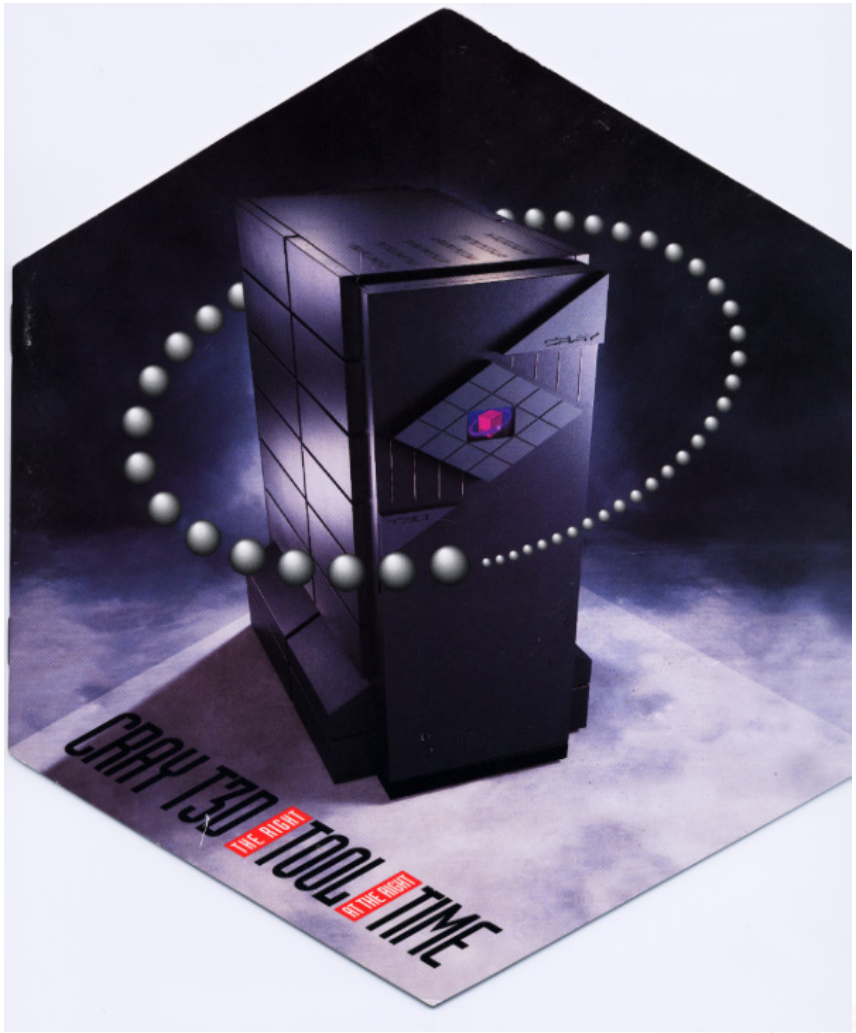

It wasn't exactly NERSC SSP, but it was the closest thing to it at the time…

NAS Parallel Benchmark Results 10-93

David H. Bailey, Eric Barszcz, Leonardo Dagum and Horst D. Simon

RNR Technical Report RNR-93-016

October 27, 1993

Abstract

The NAS Parallel Benchmarks have been developed at NASA Ames Research Center to study the performance of parallel supercomputers. The eight benchmark problems are specified in a "pencil and paper" fashion. In other words, the complete details of the problem to be solved are given in a technical document, and except for a few restrictions, benchmarkers are free to select the language constructs and implementation techniques best suited for a particular system.

This paper presents performance results of various systems using the NAS Parallel Benchmarks. These results represent the best results that have been reported to us for the specific systems listed.

NAS Parallel Benchmarks on the T3D…

- At Cray, we implemented the entire NPB parallel suite with exactly two communication calls

```
CALL GET(SRC, DEST, NWORDS, STRIDE, PE)

CALL BARRIER()
```

- Bob Numrich came up with this idea while we were waiting for the development guys to get PVM to run…
- This went on to become the "SHMEM" stuff we know today
- Bob followed this idea up with F– which eventually became Co-Array FORTRAN

A Few Results From October, 1993
BT Simulated CFD Benchmark

System	Date	N Proc	Performance (Y-MP 1 Units)
Cray Y-MP	Aug 1992	1	1.0
		8	6.95
BBC TC2000	Dec 1991	112	0.58
Thinking Machines CM-2	Dec 1991	64K	2.14
Thinking Machines CM-5	May 1993	128	6.66
Intel Paragon	Sept 1993	256	8.12
Intel Paragon	Sept 1993	64	5.19
Cray T3D	May 1993	64	5.16
		128	10.04

NAS Parallel Benchmark Results 10-93

David H. Bailey, Eric Barszcz, Leonardo Dagum and Horst D. Simon

RNR Technical Report RNR-93-016

October 27, 1993

Vector's Days are Numbered…

Benchmark Tests on the New IBM RISC System/6000 590 Workstation

HARVEY J. WASSERMAN

Computer Research Group, Los Alamos National Laboratory, Los Alamos, NM 87545

Published in 1995

ABSTRACT

The results of benchmark tests on the superscalar IBM RISC System/6000 Model 590 are presented. A set of well-characterized Fortran benchmarks spanning a range of computational characteristics was used for the study. The data from the 590 system are compared with those from a single-processor CRAY C90 system as well as with other microprocessor-based systems, such as the Digital Equipment Corporation AXP 3000/500X and the Hewlett-Packard HP/735.

1 INTRODUCTION

The IBM RISC System/6000, first introduced in 1990, was an important step in the development of high-performance microprocessor systems. Using a *superscalar* architecture, the RS/6000 achieved performance on floating point intensive

2 RISC SYSTEM/6000 MODEL 590 ARCHITECTURE

For brevity, model 590 is referred to as the RIOS-2 and the older system as the RIOS-1. *Superscalar* means that the RIOS-1 processor is capable of issuing four different instructions: a branch, a

Making Vectors Cheaper - Cray J90 – 1996

- Completely compatible with the Cray Y-MP
- Provided interim service as the C90 is moved to Berkeley
- Followed later by a few Cray SV1 systems
- CMOS-based systems to reduce cost
- Leveraged "UNICOS", compilers and so on
- Last systems at NERSC that ran the Cray Floating Point Format (vs. IEEE)

Major Innovations
- CMOS Implementation
- Reduced Cost
- 32 Processor SMP

Mcurie: Cray T3E MPP - 1996

NERSC transitions users to MPPs

- First 128 processor has factory trial in August of 1996
- Installation at Berkeley Lab happens in September
- A second 512 processor T3E is installed and accepted in 1997
- The Checkpoint Restart Finale
- Production Scheduling
- Fun Stories
  - The R5 debacle
  - 2 Bills & 4 Steves

Major Innovations
- E-Registers
- Stream Buffers
- UNICOS mk
- Co-Array FORTRAN
- Checkpoint / Restart

1998 Gordon Bell Prize Winner – First Teraflop Application – We Were Proud of This Group…

BERKELEY LAB | SCIENCE ARTICLES ARCHIVE

NERSC Scientist Shares Supercomputing's Top Prize

November 20, 1998

By Jon Bashor, jbashor@lbl.gov

Andrew Canning, a member of NERSC's Scientific Computing Group, was part of an international team that won the 1998 Gordon Bell Prize for the best achievement in high-performance computing. The winners of the prestigious award were announced during SC98, an annual conference on high-performance computing and networking held in Orlando, Florida. Canning's collaborators include scientists at Oak Ridge National Laboratory (ORNL), the Pittsburgh Supercomputing Center and University of Bristol (UK).

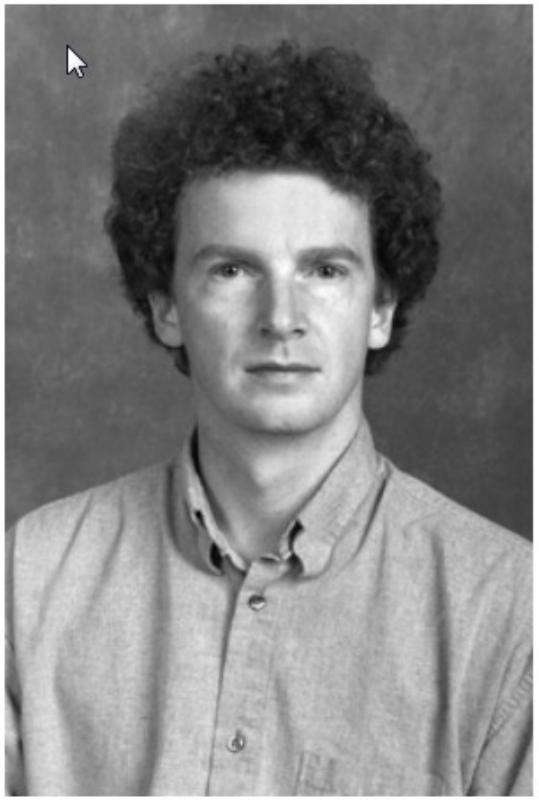

Andrew Canning

This Happened Along The Way

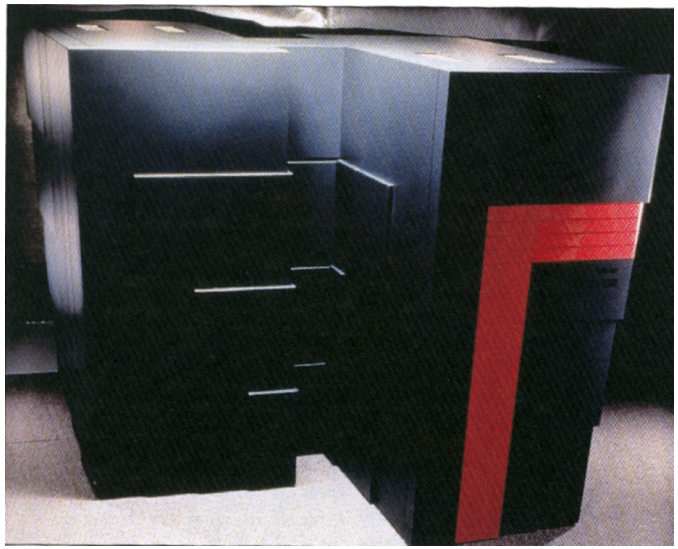

This Wasn't a Cray

But it Turns Out it Was Sold to NERSC by a Cray Guy…

Franklin: Cray XT4 - 2007

- AMD Opteron
- Brief pit-stop with air-cooling
- Largest Cray XT4 installed
- Upgraded from dual-core to quad-core
  - Quadrupled peak performance
- Early adopter of Cray Linux Environment
  - Earlier Systems were based on Catamount
- DVS developed to allow connectivity to NERSC global File System

Major Innovations
- Seastar 2
- CLE
- DVS

It Was Really a 108 Cabinet System (but 6 were empty, thanks to the pole…)

(fineartamerica.com/products/cray-xt4-supercomputer-cluster-lawrence-berkeley-national-laboratory-ca

How Do We Name These Things?

- We needed a name for the “productized” version of Red Storm
- Cray RS ?
- William White had these plastic, injection-molded letters on his door from the earlier Alpha cluster project..

CRAY T3X

Hopper: Cray XE6 2011

- Started Life as an interim XT5
- Gemini and SeaStar were made plug compatible
- 24-core node
- Warm Swap
- Much less "fragile" than Cray XT4
- Much better MPI latency

Major Innovations
- Gemini Interconnect
- Warm Swap
- Ecophlex
- Cool Murals…

How Do We Name These Things?

ECOPHLEX ???

ECO

PHase

Liquid

EXchange

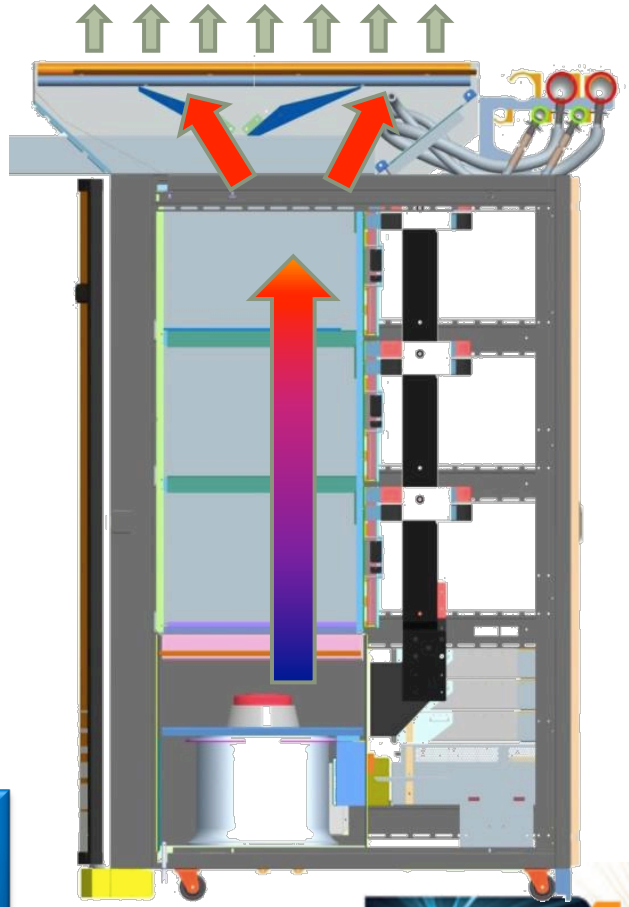

Rejected Marketing Slogan: "You know it's powerful by it's mighty roar !!!"

Edison: Cray XC30 - 2011

- First XC30 installed at a customer site
- Another "interim" plan
- Ivy Bridge change in plan…
- X86 again, but different
  - Gemini -> Aries
  - AMD -> Intel
  - Copper -> Optics
  - Vertical -> Horizontal
  - Loud -> Quiet
- We tried to make this a smooth transition for users and administrators
  - Same CLE
  - Same Compilers & Libraries.

Major Innovations
- Aries Adaptive Routing
- Free Transverse Cooling
- Dragonfly

What? Edison was a used machine ?

Changed into new clothes and delivered to Oakland in late 2012

Cray announced the XC30 system at SC12

Scalability – NERSC "Now Computing" Snapshot (taken Sept. 4th 2013)

A small sample of jobs now running on NERSC's supercomputers. [Full Screen Display]

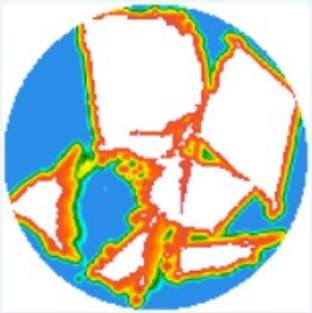

Chombo-Crunch: Advanced Simulation of Subsurface Flow and Reactive Transport Processes Associated with Carbon Sequestration (m1792)

DOE Office	Advanced Scientific Computing Research	**Science Area**	Geoscience	**Investigator**	David Trebotich, Lawrence Berkeley National Laboratory
Computer	Edison	**Cores**	65,544	**Core Hours Used**	223,981.7

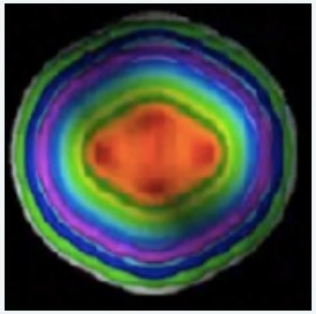

Static and dynamic solutions for heavy nuclei (m1451)

DOE Office	Nuclear Physics	**Science Area**	Nuclear Physics	**Investigator**	Ionel Stetcu, Los Alamos National Laboratory
Computer	Hopper	**Cores**	28,968	**Core Hours Used**	33,910.1

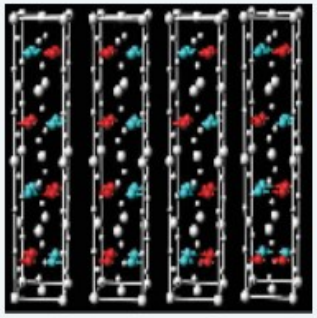

Spin-lattice Coupling in Magnetic Phase Transition (m891)

DOE Office	Basic Energy Sciences	**Science Area**	Materials Science	**Investigator**	Yi Wang, Pennsylvania State University
Computer	Carver	**Cores**	256	**Core Hours Used**	903.3

Edison

Hopper

Carver

Its Been a Great 40 Years !

CRAY

We build computational tools that help change the world!

NERSC

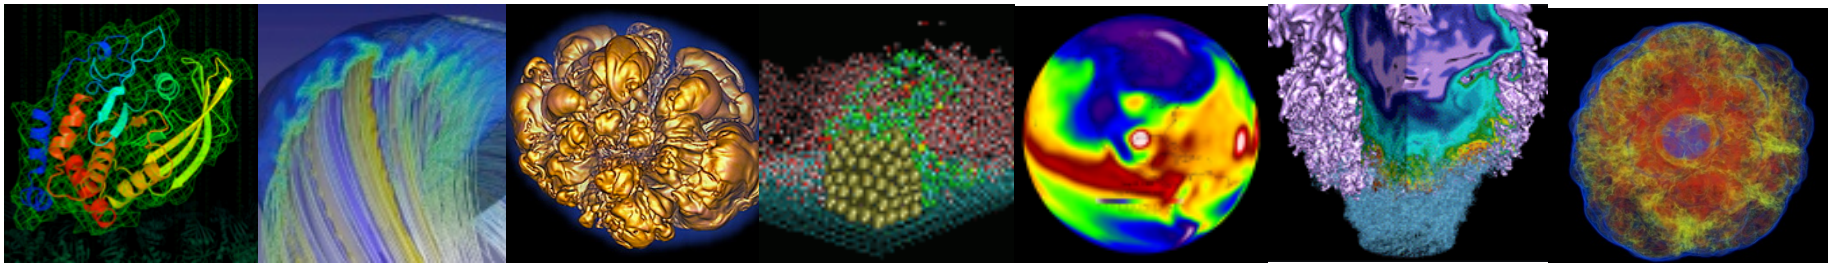

We know it's our customers who do the real heavy lifting !

NERSC 40 YEARS at the FOREFRONT 1974-2014

Thanks For a Great 40 Years!