Trends in Supercomputing and Evolution of the San Diego Supercomputer Center:

Jan 11, 2017

Richard Moore

Recently retired as UCSD/SDSC Deputy Director

Outline

- Introduce myself
- What puts the super in supercomputing?
- Trends in supercomputing
- Evolution of San Diego Supercomputer Center

My background

- **Undergrad at U Michigan (Applied Math, Astronomy)**
  - **And wine, women and jazz**
- **PhD at U Arizona (Optical Astronomy)**
  - **Optical polarization of quasars, and unusual features of the small class of highly-polarized quasars and BL Lacertae objects**
  - **=> Quasars are driven by black-hole accretion; relativistic jets likely to cause extreme characteristics in a subset of quasars**
- **Postdoc at Caltech (Radio Astronomy)**
  - Optical polarimetry at Palomar 200”
  - **Deployed K-band maser receiver (JPL) at Owens Valley Radio Obs**
  - **Very Long Baseline Interferometry (VLBI) imaging**
  - **Superluminal acceleration (relativistic jets)**

First career: Observational astronomer

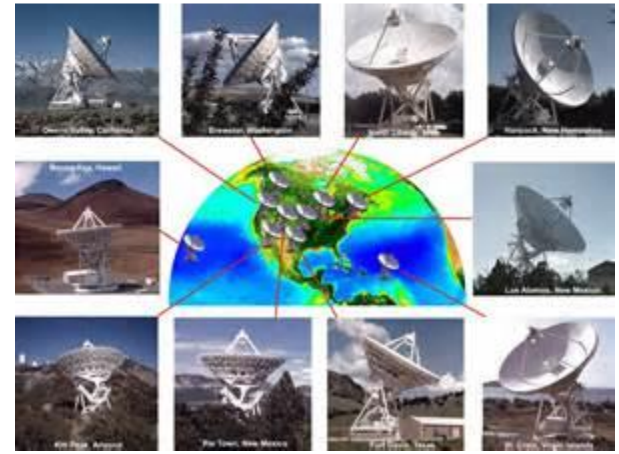

Second career: Aerospace

- **Aerospace Corporation, El Segundo**
  - **FFRDC for Air Force Space Systems LA**
- **Photon Research Associates, San Diego (early SAIC spin-off, now Raytheon division)**
  - **Physics-based modeling/simulation of targets and backgrounds (primarily IR, some radar and visible), sensors and digital processing**
  - **Missile warning systems, ballistic missile defense, other classified systems**

Third career: Supercomputing

- **Joined SDSC in 2002**
- **Held various leadership roles thru several transitions**
  - **Final steps of GA -> UCSD transition**
  - NSF's disaggregation of center **funding to individual competitions**
  - **Increased engagement w/ UCSD and UC**
- **Retired mid-2015**

Outline

- Introduce myself
- What puts the super in supercomputing?
- Trends in supercomputing
- Evolution of San Diego Supercomputer Center

What puts the super in supercomputing?

- **Ability to apply large computational capability to a single problem**
- **Supercomputers are not just large computing installations (capacity)**
  - **A few internet companies far dominate the computing landscape in terms of** capacity; ~25% of Intel's chips go to four web companies.
- **You can always do capacity computing on a capability machine, but *you can't do capability computing on a capacity machine***
- **In early Cray/CDC days, the super part was very fast processors and vector processing**
- **But clock frequencies maxed out (energy limits) and economics now drives supercomputing to rely largely on commodity components**

Parallelism is now the super in supercomputing

- **Supercomputers have high bandwidth/low latency interconnects that pass data between nodes so that many nodes work in parallel on same problem**
  - **Current interconnects: ~+100 Gbps bandwidth, ~100 ns latency per hop**
  - **Interconnect network adds ~10-30% of cost of system**
  - **Software protocols for passing data between nodes**
- **Common interconnects**
  - Ethernet (latencies relatively high … standard for ‘web farms’)
  - **InfiniBand (open standard, but only vendor left is Mellanox)**
  - **Intel OmniPath** – recent ‘interconnect in the chip set’
  - **Proprietary networks by Cray, IBM, Japanese and Chinese**
- **Nodes connected via node interface cards and switches**
  - Scaling switches to large numbers of ‘ports’ is an N^2 **challenge**
  - **Largest switches (IB EDR or 100G Ethernet) ~500 ports and a LOT of bandwidth in a single box (100+ Tbps).**
  - **Can build multi-level switch architectures, but more hops (i.e. latency) and more ports for inter-level connections**

Hardware is just the beginning … making it work efficiently is harder

- **Computational efficiency = flops achieved / theoretical peak flops**
- **How much time is spent moving data vs doing calculations?**
  - **Memory bandwidth/latency are key – and are lagging behind flops**
  - **Floating-point registers often wait for the data to operate on**
- **'Standard' benchmark for HPC systems ~50-80% efficiency**
- **But many scientific applications have single-digit efficiencies**
  - **Scaling algorithms up to larger numbers of cores/nodes is hard, and memory access**
- **Need to decompose scientific problems and develop algorithms that map well onto core/node/interconnect/storage architecture**
  - **Architectures change frequently**
  - 'Co-design' is the buzzword to design algorithms/architecture together

Outline

- Introduce myself
- What puts the super in supercomputing?
- Trends in supercomputing
- Evolution of San Diego Supercomputer Center

Emerging processor technologies

- **Intel dominates x86 CPU chips and keeps increasing cores/chip (also AMD, IBM Power, etc.)**
  - **Balance with memory bandwidth is the challenge**
- **“Accelerators” (high flops/$ and flops/watt)**
  - **Intel Phi co-processor (~64 cores, 4 threads/core)**
  - **Graphics processors (e.g. NVIDIA)**
  - **Field-programmable gate arrays (FPGAs)**
  - **Memory quantity/access/speed is a challenge with most accelerators**
- **ARM processors?**
  - **Not yet a major player, but strong power and economics drivers**

System 'balance' changing dramatically: Flops are cheap, especially w/ multi-core chips and accelerators

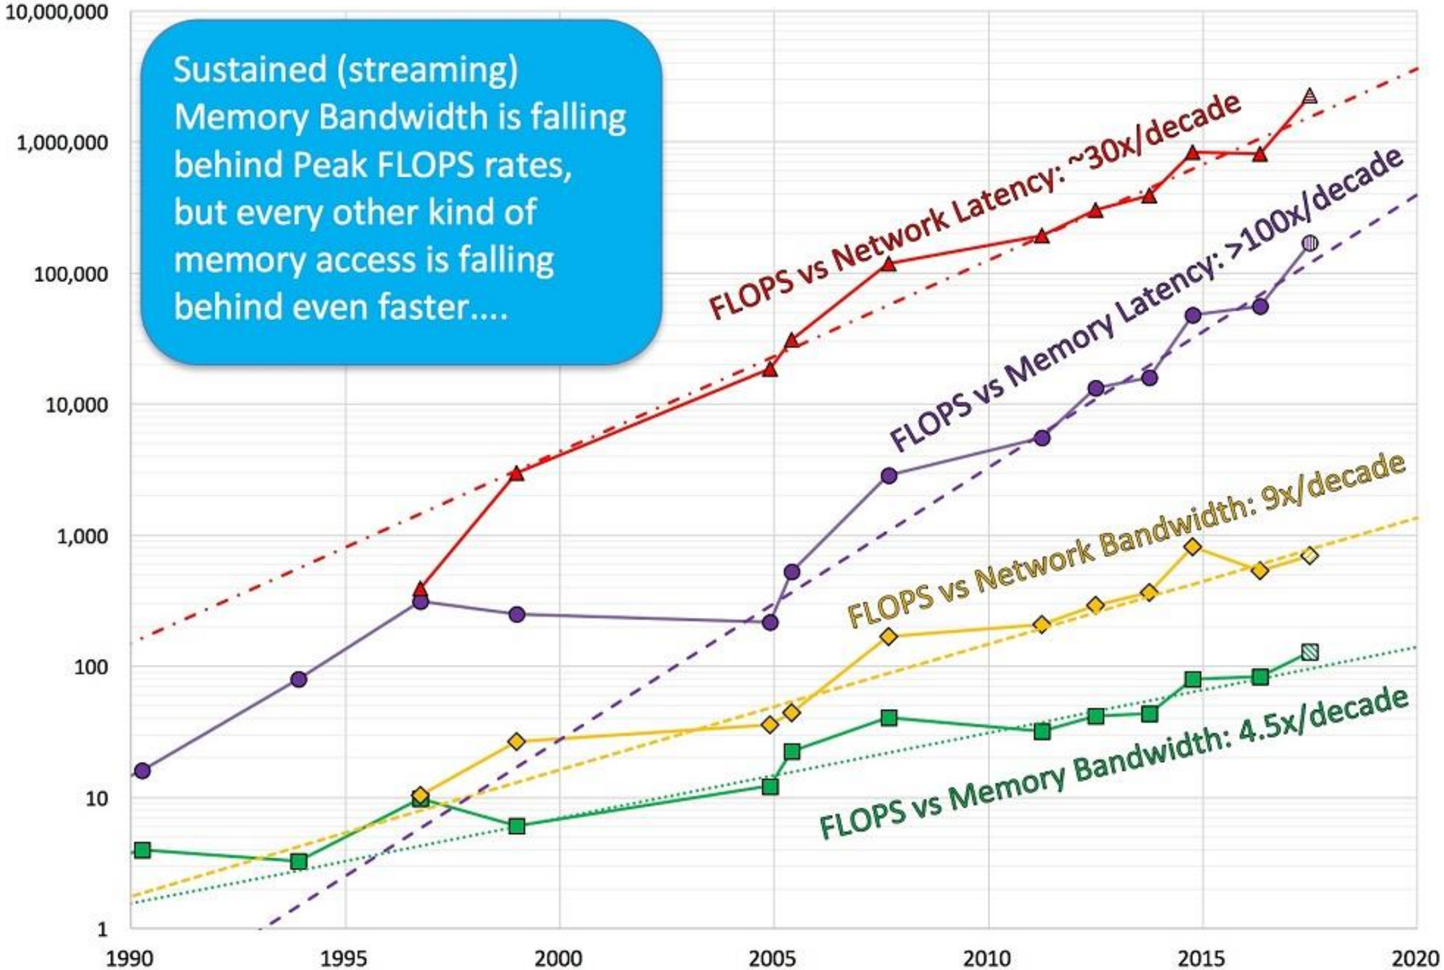

From John Macalpin, UT/TACC

For really large systems (e.g. exascale), other issues become limiting factors

- **System reliability - can all components work long enough to complete the calculation?**
  - **Fault tolerance against many forms of errors**
  - **Checkpointing (copying system state) – at reasonable speeds**
- **Energy consumption**
  - **Largest systems now >10MW**
  - DOE set ~20MW requirement for '**exascale**' system (~$20**-50M/yr utilities)**
- **Fewer and fewer applications are making effective use of entire system**

Current "Top 10" Supercomputers (www.top500.org, Nov 2016)

Organization	Name	Peak (PF)	Benchmark (PF)	Vendor/ Network/ Accelerator	Cores (million)	Power (MW)
National Supercomputing Center Wuxi	TiahuLight	125.4	93.0	NRCPC/Sunway/Sunway	10.65	15.4
National Supercomputer Center Guangzhou	Tianhe 2	54.9	33.9	NUDT/TH-Express/ Phi	3.12	17.8
Oak Ridge	Titan	27.1	17.6	Cray/Gemini/NVIDIA	0.56	8.3
Livermore	Sequoia	20.1	17.1	IBM/BlueGene/ None	1.56	7.9
NERSC	Cori	27.9	14.0	Cray/Aries/Phi	0.62	3.9
JCAHPC (Japan)	Oakforest	24.9	13.6	Fujitsu/OmniPath/Phi	0.56	2.7
RIKEN (Japan)	'K'	11.3	10.5	Fujitsu/Tofu/None	0.71	12.7
CSCS (Switzerland)	Piz Daint	16.0	9.8	Cray/Aries/NVIDIA	0.21	1.3
Argonne	Mira	10.1	8.6	IBM/BlueGene/None	0.79	3.9
Sandia	Trinity	11.1	8.1	Cray/Aries/None	0.30	4.2

Supercomputer 'Top500' Performance over time (www.top500.org)

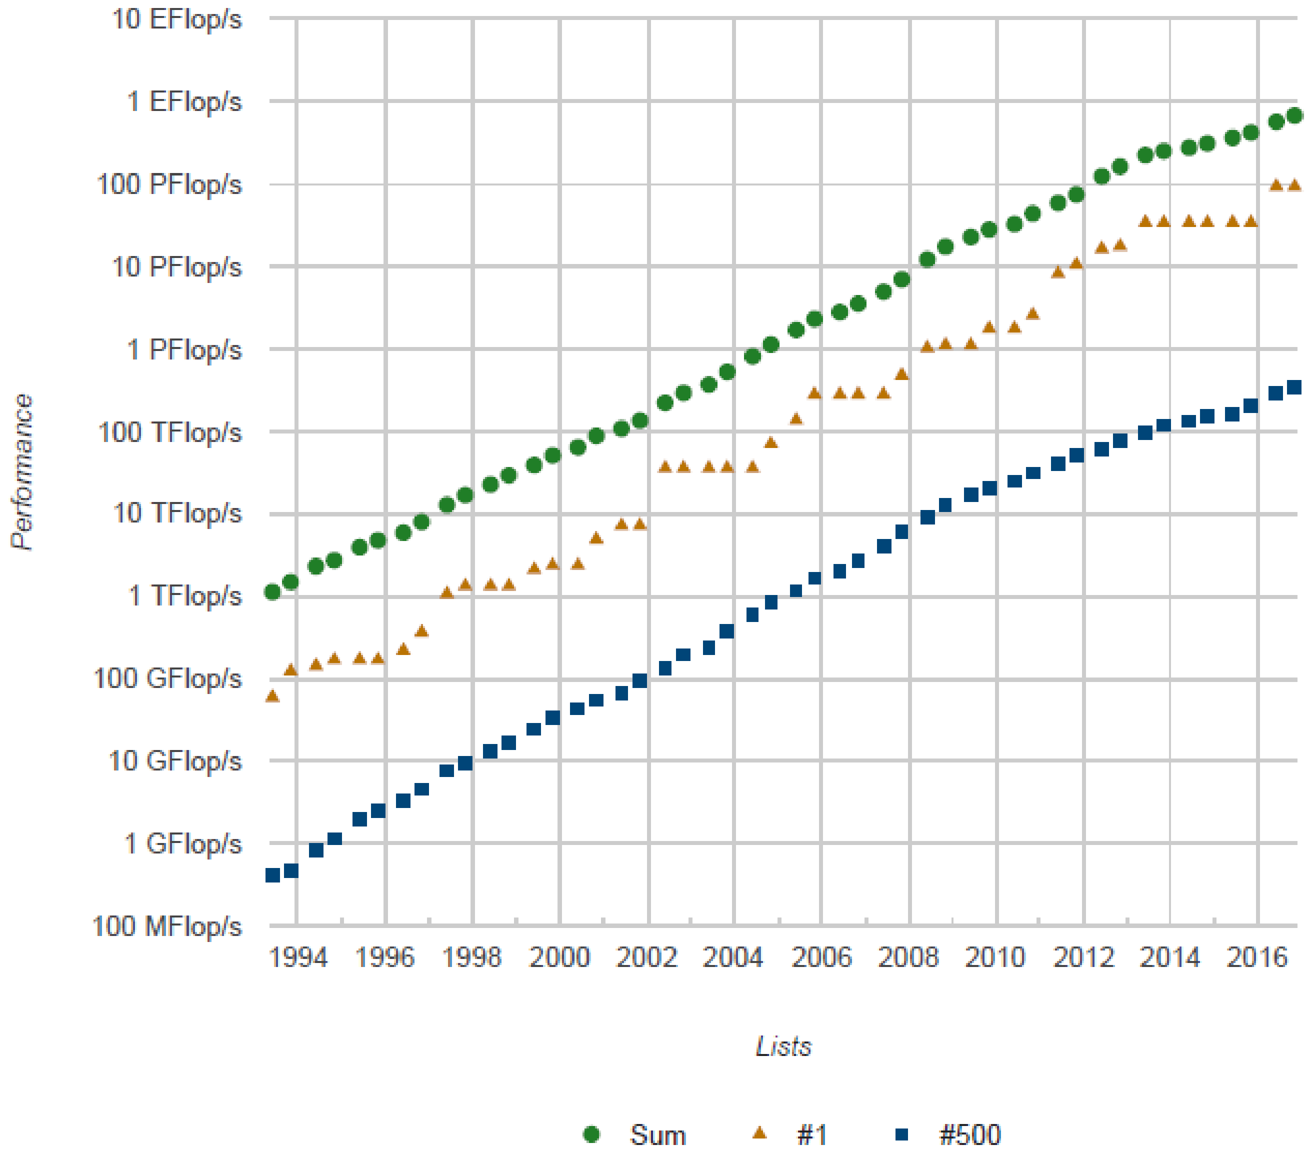

SDSC

UCSD

Top 500 trends

- **Accelerator/CP family – performance/systems**
  - **Intel and NVIDIA are battling it out**
- **Vendors – performance/systems**
  - **Cray is major player in large systems,**
  - **IBM fading (especially with end of BlueGene series)**
  - **HP has many of the smaller systems,**
  - … many others
- **Countries – performance/systems**
  - **China!**

Accelerator family (2006-2016) – Systems/Performance Share (from top500.org)

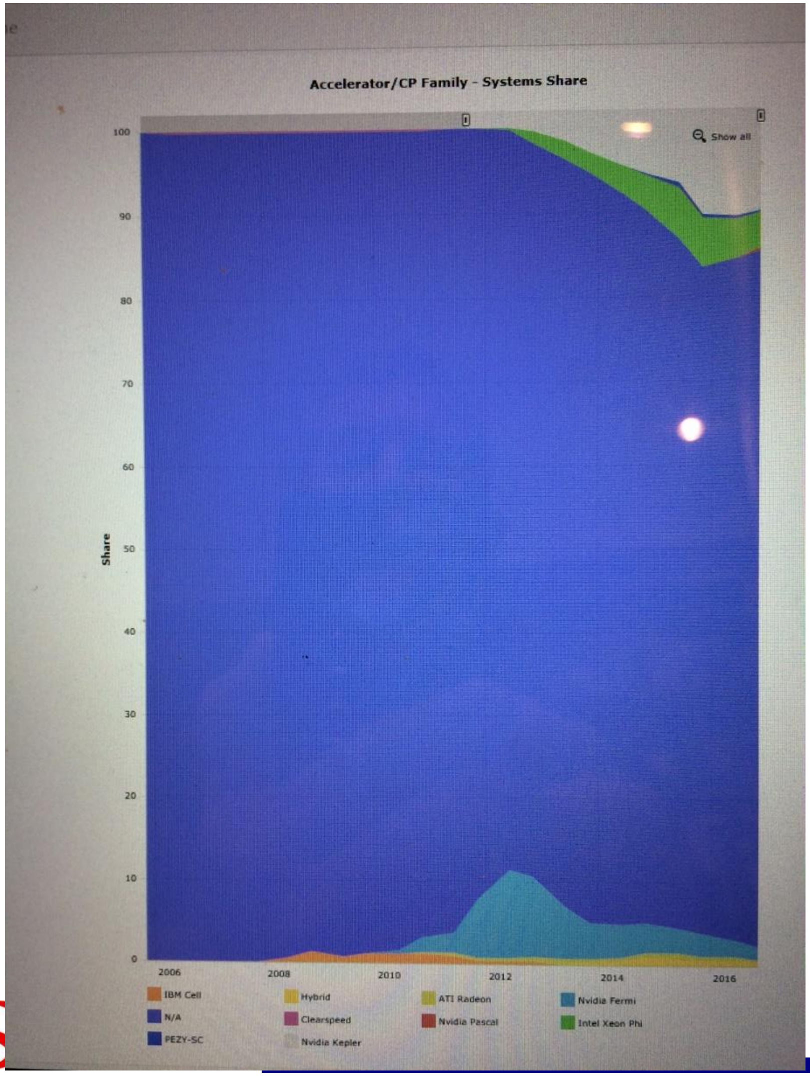

Vendors (1996-2016) – Systems/Performance Share (from top500.org)

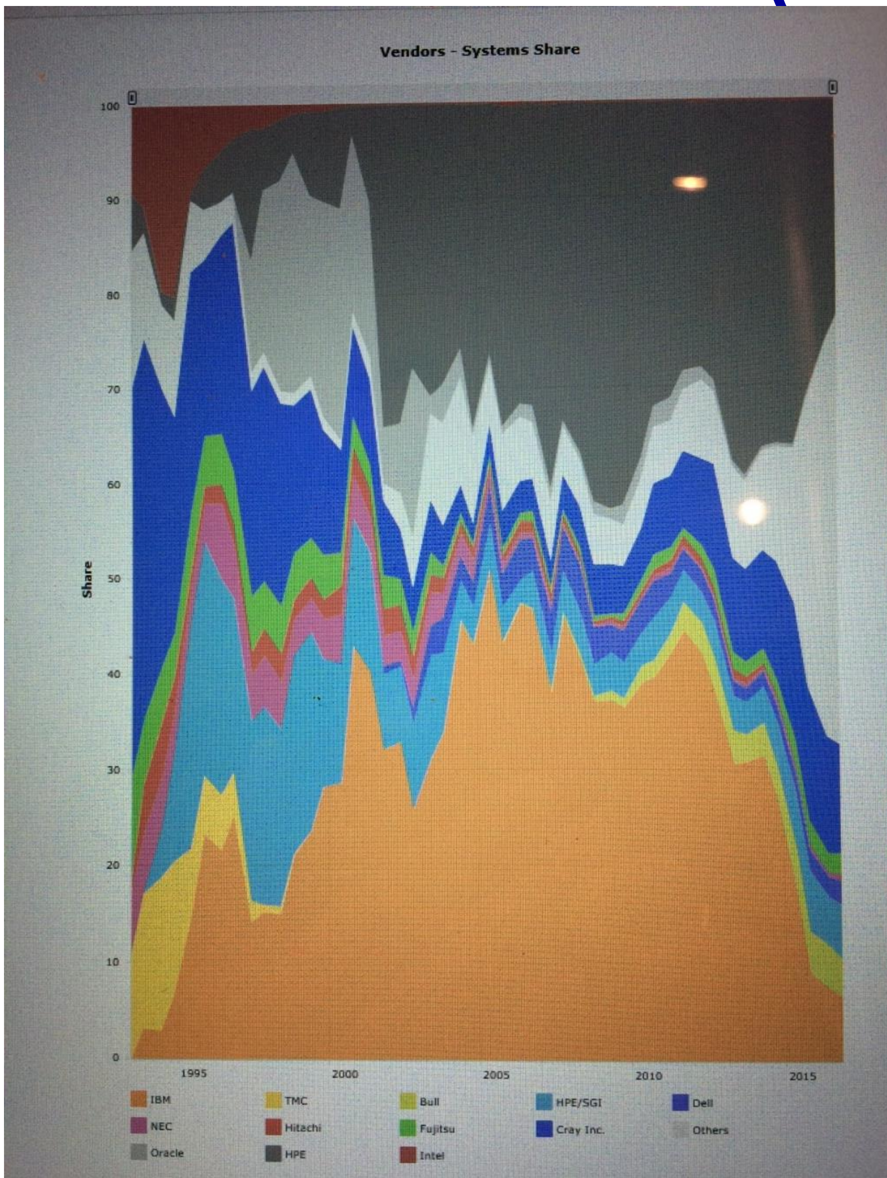

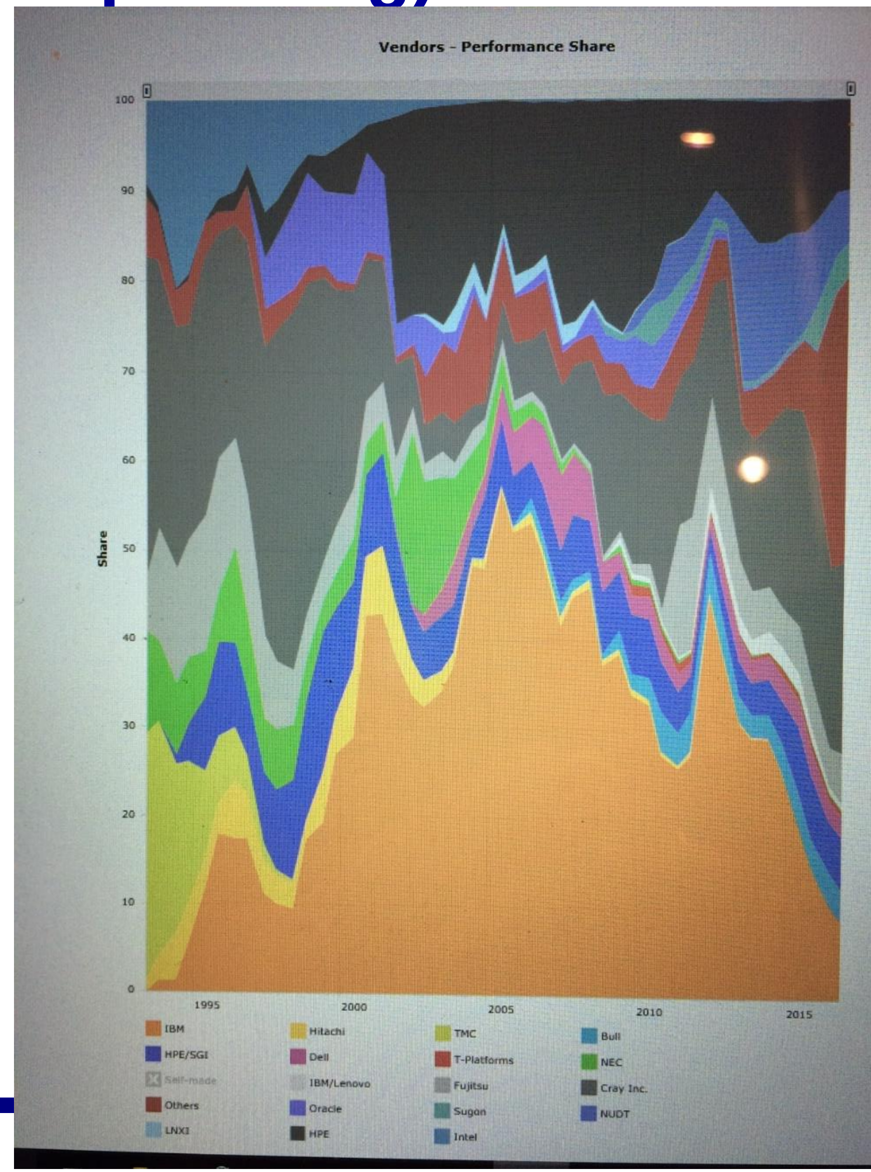

Countries (1994-2016) – Systems/Performance Share (from top500.org)

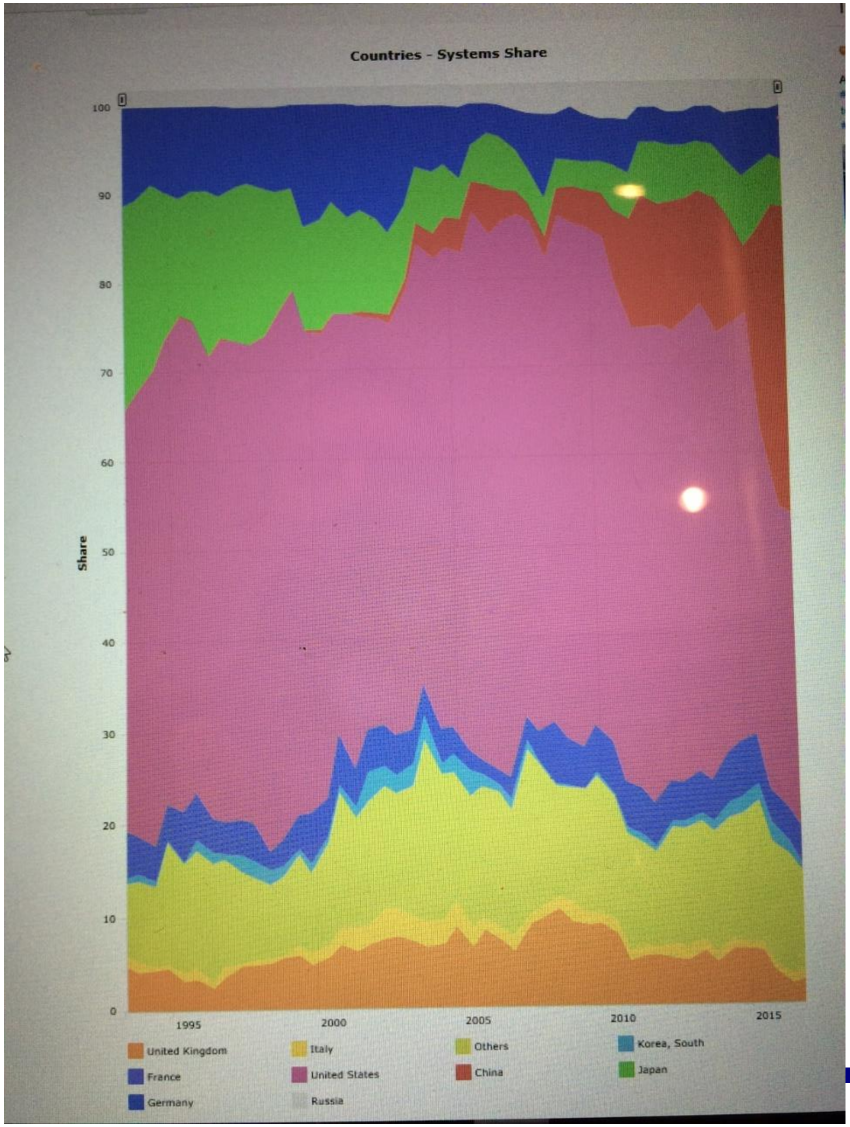

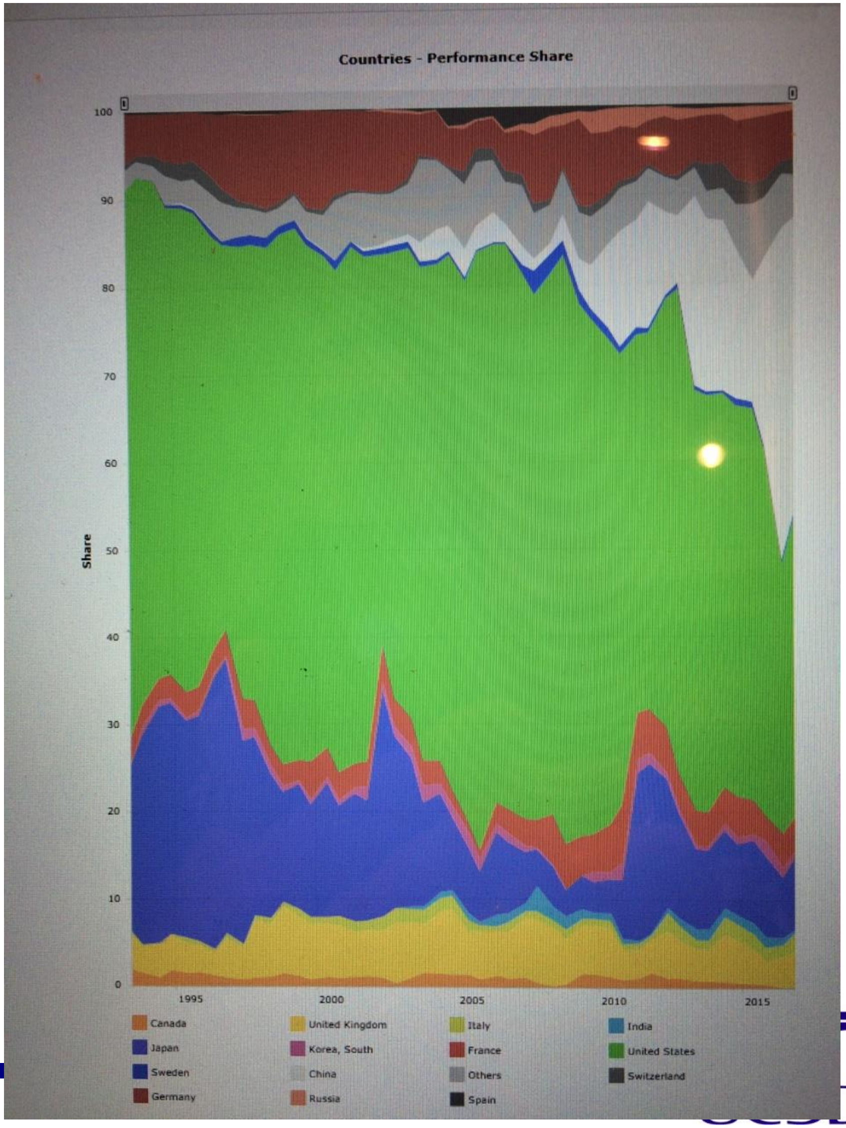

Outline

- Introduce myself
- What puts the super in supercomputing?
- Trends in supercomputing
- Evolution of San Diego Supercomputer Center

SDSC – A Pioneer in Supercomputing

- **Established as a national supercomputer resource center in 1985 by NSF**
  - **Run by General Atomics – rare for NSF**
- **Became an Organized Research Unit of UC San Diego in 1997**
  - **Second largest ORU at UCSD**
- **Grant revenues ~$20-60M/yr (~$25M/yr)**
  - **Full-time employees 100-300 (now ~220)**
- **>$1B revenue over lifetime!**
- **World leader in data-intensive computing and data management**

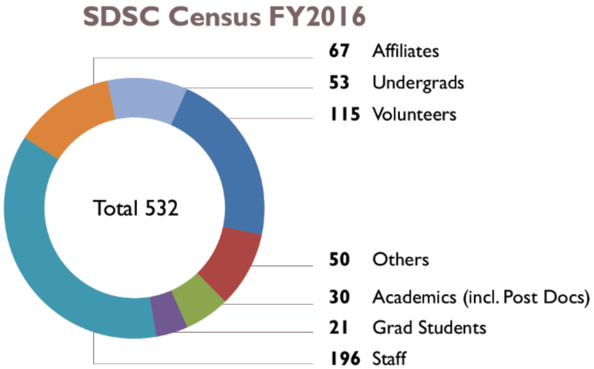

SDSC

Evolutionary Phases of SDSC

- **1985-1997: NSF national supercomputer center; managed by General Atomics**

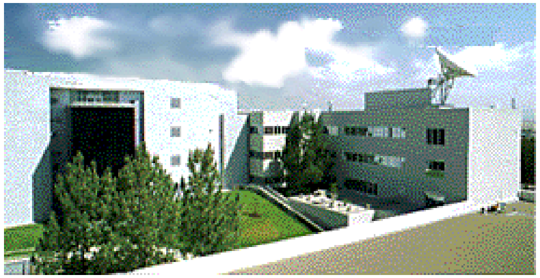

- **1997-2007: NSF PACI program leadership center; managed by UCSD**
  - **SDSC led National Partnership for Advanced Computational Infrastructure w/ ~35 universities – supercomputers, technology and applications**

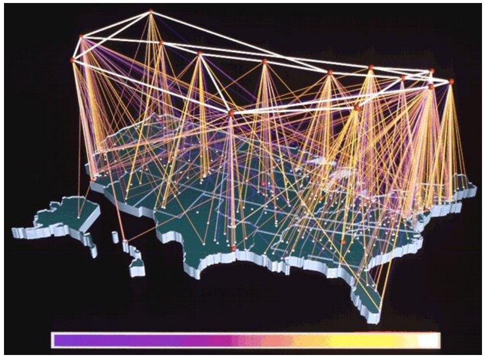

- **2008-future: multi-constituency cyberinfrastructure (CI) center**
  - Still NSF national "resource provider"
  - **But provides data-intensive CI resources, services, and expertise for campus, state, and nation**

In the beginning there was Sid … submitting an unsolicited proposal to NSF

- **“proposal for the development and operation of a highly interactive, user-oriented supercomputer facility”**
- **Cray X-MP + VAX-based user centers at UCSD, SIO, Salk, SDSU, ….**
- **User training and consulting**

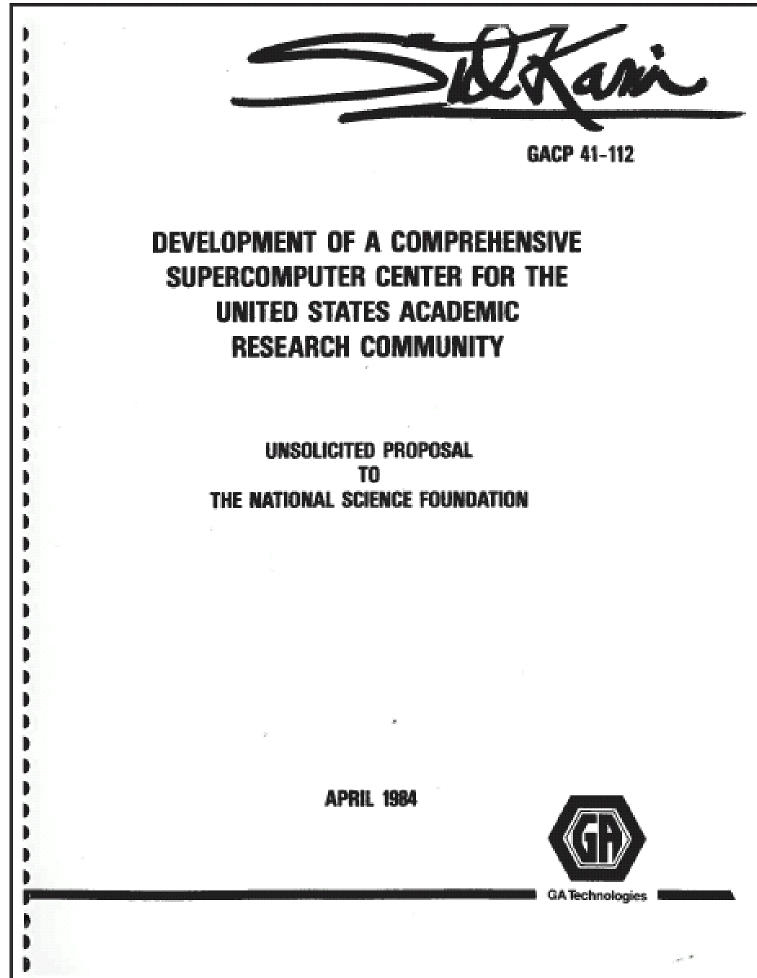

GACP 41-112

DEVELOPMENT OF A COMPREHENSIVE SUPERCOMPUTER CENTER FOR THE UNITED STATES ACADEMIC RESEARCH COMMUNITY

UNSOLICITED PROPOSAL
TO
THE NATIONAL SCIENCE FOUNDATION

APRIL 1984

GA Technologies

A Cavalcade of Supercomputers

1985 – Cray X-MP

1989 – Cray Y-MP

1993 – Cray C90

1996 – Cray T3E

architectural convergence

1990 – NCUBE 2

1993 – Intel Paragon

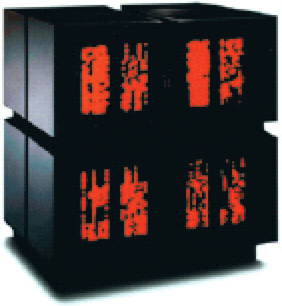

1994 – Thinking Machines CM-2

National Partnership for Advanced Computational Infrastructure (NPACI)

SDSC * U Texas * U. Michigan * Caltech + ~30 other partners

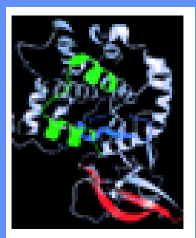

Molecular science

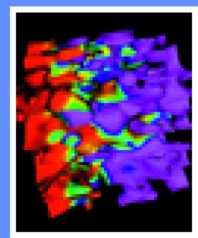

Engineering

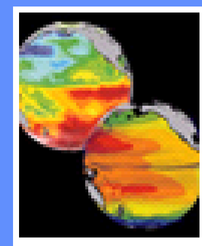

Earth Systems

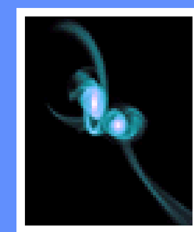

Astrophysics

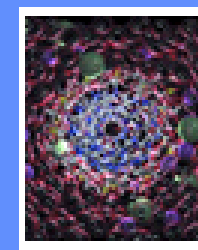

Bioinformatics

applications

NPACI

technologies

Data management

SCALABLE VISUALIZATION

Scientific visualization

NPACIgrid

Grid computing

ROCKS

www.rocksclusters.org

Cluster management

2000: First Academic TeraFlop* Supercomputer IBM Blue Horizon at SDSC

***1 trillion floating point operations per second**

2004: Data Star
Targeting Data-Intensive HPC Applications

IBM Power4 Data Star
2000 cpu, 17 Tflops
GPFS filesystem

- **TeraShake virtual earthquake simulation**

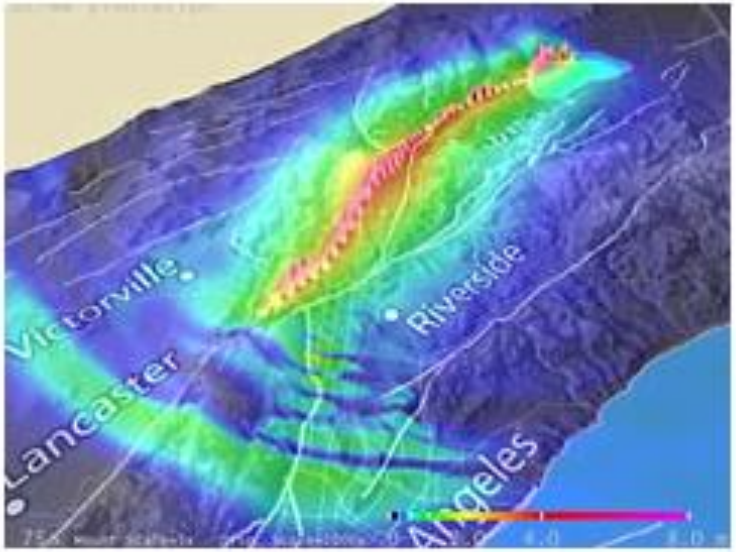

Amit Chourasia, SDSC Visualization Services

SDSC leveraged its expertise and expanded to include many PI-led labs

Next Generation Biology Workbench

Groups & Labs

- Advanced Cyberinfrastructure Development Group
- Advanced Query Processing Laboratory
- Applied Network Research
- Bourne Laboratory
- Complex Systems
- Cryoelectron Microscopy and 3D Image Reconstruction
- Cooperative Association for Internet Data Analysis
- Homeland Security
- Laboratory for Computational Astrophysics
- Laboratory for Environmental and Earth Science
- Molecular Interaction and Crystallography
- Next Generation Tools for Biology
- Oceanography and Biodiversity Research Group
- Pacific Rim Activities
- Performance Modeling and Characterization Laboratory
- Scientific Workflow Automation Technologies Laboratory
- SDSC Education Group
- Spatial Information Systems Laboratory

NSF's Big Experiment with 'Grid Computing' (2002 – ~2006)

- **SDSC founding member of TeraGrid, together with NCSA, Caltech, and Argonne Nat'l Lab**
  - **Four HPC sites interconnected @ 40 Gbps (a lot in 2002!)**
- **Original 'grid' concept evolved to what is now an interconnected collaboration of resource providers**

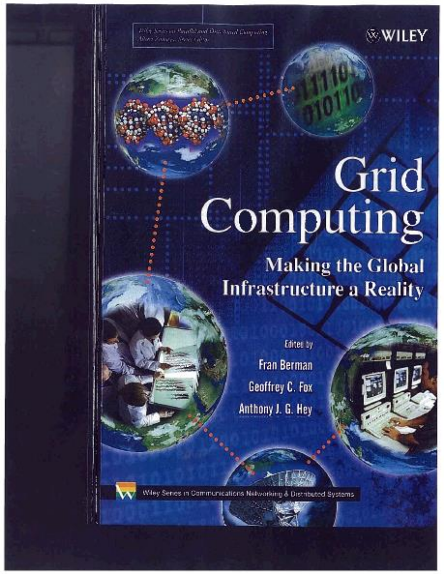

Third Phase for SDSC

- Change of NSF funding model away from 'Center' model led to a re-visioning of SDSC
- Continue role as NSF national 'resource provider' - but shorter individually-competed awards
- Continue PI-led labs and projects, with emphasis on data science
- Greater emphasis on integrating with UCSD/UC to provide resources and services of value to local researchers

SDSC building dedication Oct. 2008

Recent/Current HPC Systems at SDSC

- **Trestles (2011-2015) 100 TF system by Appro/Cray**
  - First NSF system designed for 'the 99%' – **users w/ modest-scale jobs that can use fast turnaround, + science gateways (easy portal-based interface)**

- **Gordon (2012-present) – 300 TF system Appro/Cray**
  - **Designed for data-intensive applications**
  - **First HPC system in the world to incorporate flash memory at-scale in the memory hierarchy (cache – RAM – flash - disk)**
- **Triton Shared Computing Cluster (2011-present)**
  - **A shared UCSD/UC resource, with PIs purchasing nodes in shared system**

- **Comet (2015-present) – 2PF system by Dell**
  - **Continues Trestles themes at larger-scale: modest- -scale jobs, capacity, turnaround, gateway-friendly**

High-*performance computing 'for the 99%'*

- Scientific impact occurs through-out the Branscomb pyramid
- Many HPC users just need 'capacity' to do their science

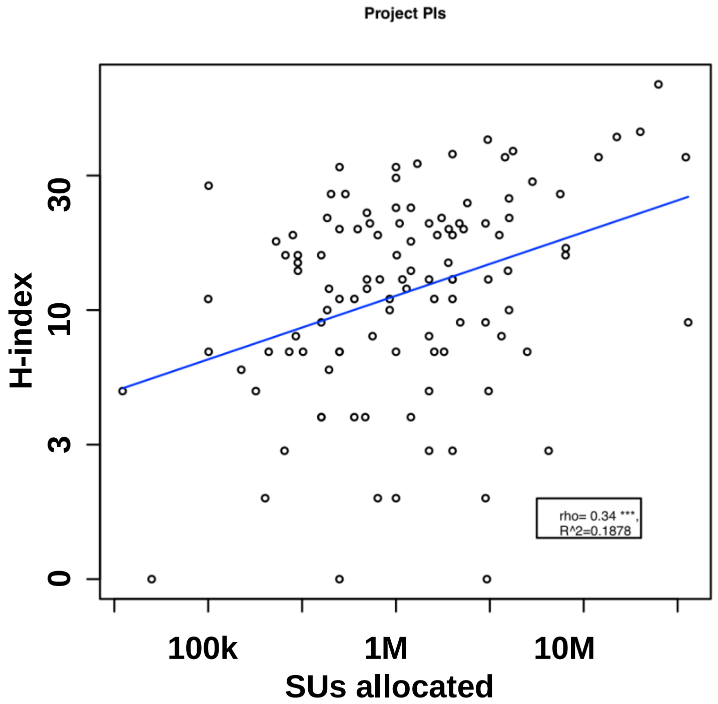

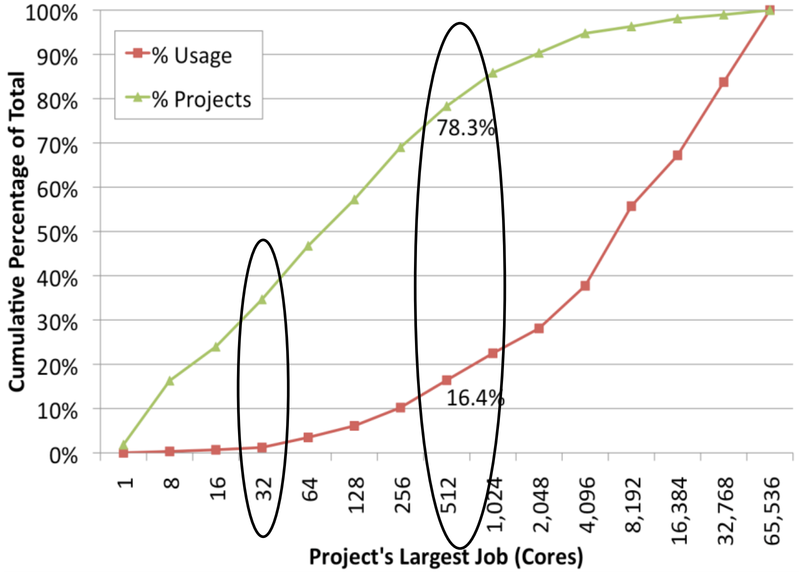

- From 2009 data across all NSF systems

Geographical Distribution of National Users of SDSC HPC Systems

- **During FY2015-16, SDSC's Gordon and Comet**
  - **Supported 2,348 unique users around the world (98% in US)**
  - **And delivered 106M and 305M core-hours respectively**

Training the Next Generation Workforce

- Have trained hundreds of local professionals in data science “boot camps” over the last 3 years
- Coursera online Big Data Specialization course has over 400,000 registrants in last year!
- Recently inaugurated the “Modern Data Science Academy” in partnership with UC San Diego Extension
- SDSC faculty participate in teaching in the Master of Advanced Studies in Data Science program

Educating Our Youth in STEM

- **StudentTECH Workshops**
  - Week long workshops offered during the summer for middle and high school students on a variety of topics, from JAVA programming, 3D printing, skateboard design to underwater robotics and artificial intelligence.
- **Research Experience for High School Students (REHS)**
  - 8 week summer internship program at SDSC. High school students from the San Diego region are selected to work with a mentor on a research project. At the end of the program, students prepare a scientific posters reflecting on their summer experience and display them at a final celebratory event.
- **Mentor Assistant Program (MAP)**
  - Campus-wide program to engage high school students (11th and 12th graders) in a mentoring relationship with an expert from a wide variety of disciplines, from history, marine science and nanoengineering to pharmacology, reproductive medicine and computer science.
- **TeacherTECH Program**
  - Professional development for K-12 teachers in computer science, data science and the use of technology. Over 1,000 teachers trained to date!

May 06, 2014 | By Jan Zverina

Teen Mentored by UC San Diego Professors Wins $250,000 in Science Prizes

Research on anti-flu drugs relies on NSF's XSEDE supercomputers

CI Services Initiative: A UC Data Center

Cold Isle Containment in New Datacenter

Hot Isle Containment in Existing Datacenter + Flooring Retrofit Phases 2&3

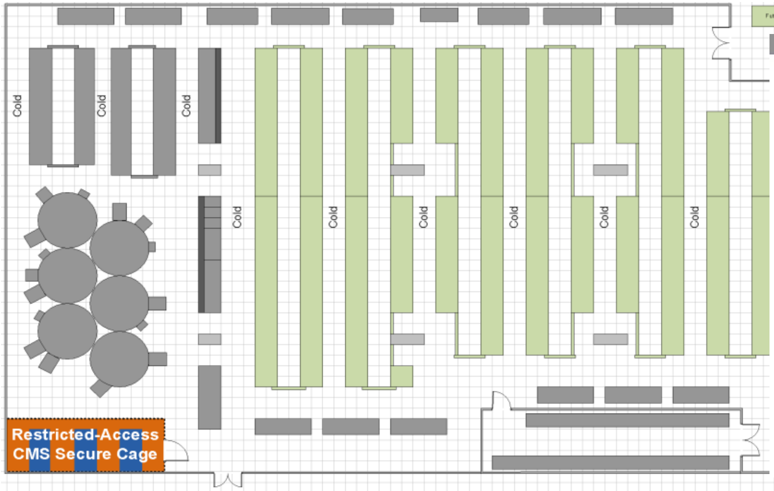

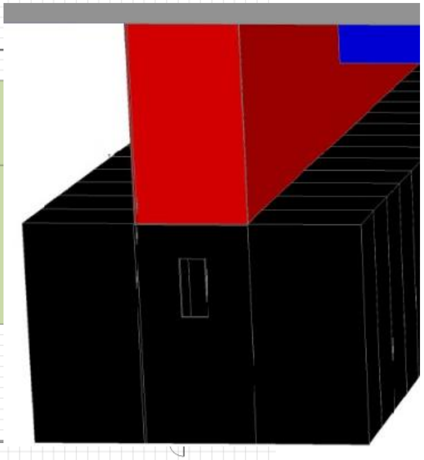

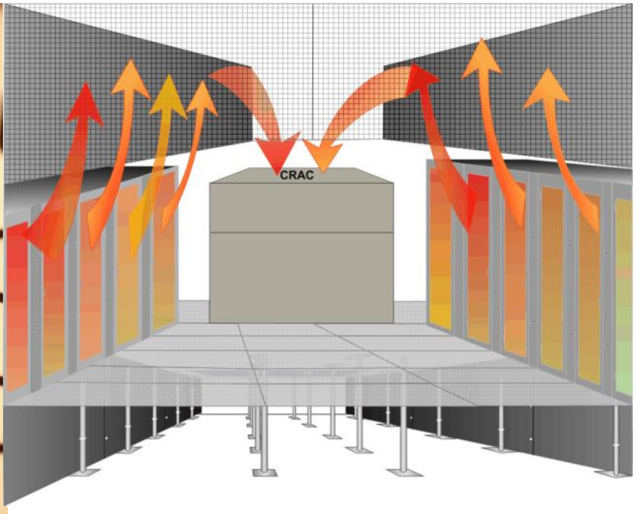

CRAC VFD Retrofits, New Sensors to Throttle with Load

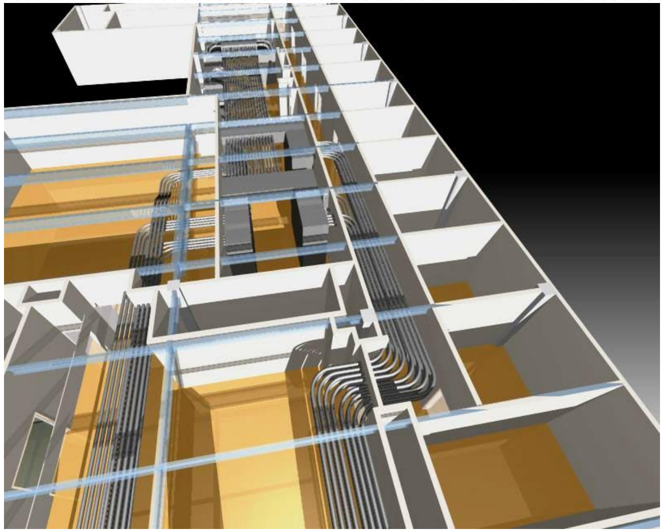

Major Power Expansion

Enabling the Future of Medicine

- **New Drug Candidate May Reduce Deficits in Parkinson's Disease**
  - UC San Diego-led team uses SDSC Resources to Displace Key Neuron-Destroying Protein
- **Promising Drug Leads Identified to Combat Heart Disease**
  - UC San Diego-Monash University Team Perform Unprecedented Sampling of Proteins
- **Underlying Molecular Networks Suggest New Targets to Combat Brain Cancer**
  - SDSC/UCSD Team Uncovers Signaling Links to Glioblastoma Factor
- **SDSC Supercomputers, CIPRES Gateway Help Define New "Tree of Life"**

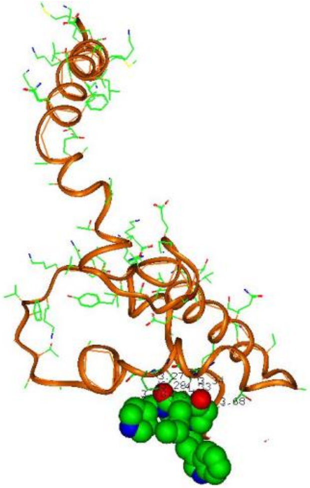

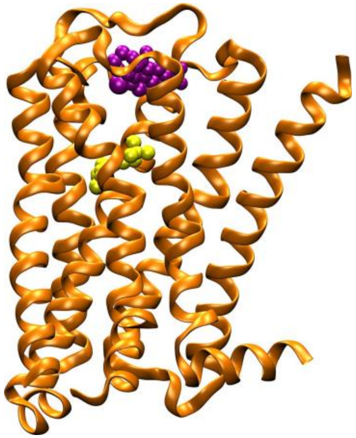

Partnering for Science & Public Safety

High Performance Wireless Research Network

- **In 2011, transitioned HPWREN from an NSF-funded science platform to a sustainable public-private partnership serving science & public safety.**
- **Played a key role in response to 2003 & 2007 wildfires**
- **Primary data link for advanced telescopes at Palomar Observatory**

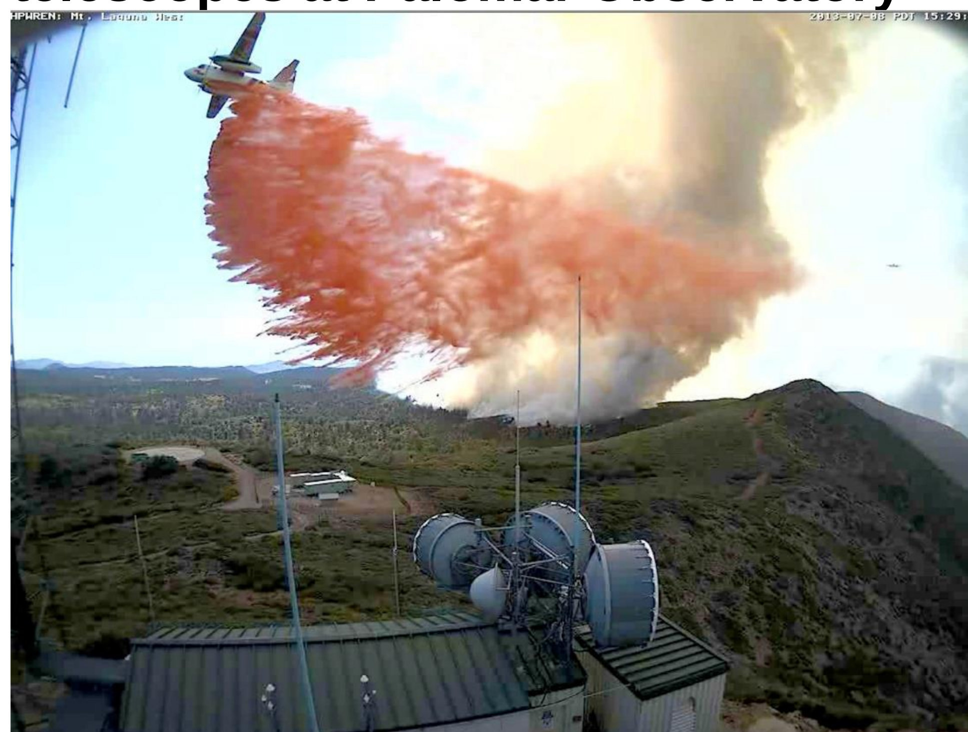

SDGE
A Sempra Energy utility

SAN DIEGO FIRE COUNTY

SAN DIEGO STATE UNIVERSITY

CAL FIRE

Thank You!