SAN DIEGO SUPERCOMPUTER CENTER at UC SAN DIEGO

SDSC

30 years

of Turning Data to Discovery

TIMELINE and TOP 30

RESEARCH and TECHNOLOGICAL ADVANCES

GA Technologies wins an NSF cooperative agreement in collaboration with UCSD to establish a supercomputer center.

SDSC opens doors; Sid Karin is SDSC's founding director.

1985

SDSC's first supercomputer, the CRAY XMP-48, arrives.

First calculations made on the SDSC CRAY on December 4, 1985 by Herbert Hamber, UC Irvine.

SCS-40 supercomputer installed at SDSC.

Over 100 peer-reviewed computational science projects authorized for NSF, Cray and SDSC allocations.

1986

SDSC, through SDSCnet, becomes first NSF center to have NSFnet access.

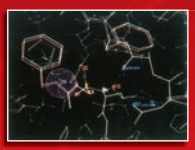

In pioneering efforts in drug design, Paul Bash, et. al., using SDSC supercomputers, determine free energies of solvation for proteins and nucleic acids, and relative free energies for binding, published in *Science*.

1987

SDSC establishes laboratory for graphics and animation.

Electronic mail gateway established at SDSC; users exchange email with peers.

TOP 30

1987

First Steps for the Rational Design of Drugs

In a pioneering effort that allowed scientists to visualize the behavior of key biological molecules in the body, a team of scientists—using supercomputing resources at SDSC—took an important first step in the relatively new arena of rational drug design. In a paper published in the journal *Science*, the team—which included Paul Bash, Peter Kollman and Robert Langridge of UCSF and U. Chandra Singh with the Scripps Clinic—reported that they had determined the relative free energies of binding for different chemical inhibitors at the same active molecular site. The result was significant since what makes one drug more powerful or effective than another is that, at the molecular level, it binds more readily at the site at which it acts.

1988

Engineering Designer Plants

It's the most abundant protein found in nature, with a whimsical name reminiscent of the company best known for its cookies. But more important than its mere abundance, RuBisCO is the key enzyme in the initiation of photosynthesis, the process by which green plants make usable energy from sunlight. In 1988, with the help of a new detector for x-ray crystallography and SDSC's CRAY X-MP, a team of scientists led by Chapman et. al. reported, for the first time, the three-dimensional structure of RuBisCO, with a subsequent goal of building a new, improved RuBisCO to engineer designer plants that would photosynthesize more efficiently, creating more food for a hungry world.

1989

Atmospheric Carbon Dioxide from Fossil Fuels

Beginning in the late 1950s, Charles Keeling from UC San Diego's Scripps Institution of Oceanography (SIO) continuously collected data on the distribution of carbon dioxide around the globe. In 1989, Keeling—with SIO colleagues Stephen Piper and Robert Bacastow, using SDSC resources—constructed a three-dimensional computer model of the terrestrial carbon cycle that took advantage of the data collected by Keeling. The model was the first to confirm the importance of fossil fuel combustion in loading the atmosphere with carbon dioxide, especially over the northern hemisphere.

Governor George Deukemejian signs legislation, introduced by Assemblyman Dominic Cortese, giving SDSC $6 million to develop state-of-the-art visualization lab at SDSC.

NSFnet backbone becomes a production network.

1988

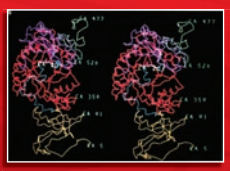

3-D structure of RuBisCo, the key enzyme in the initiation of photosynthesis and most abundant protein found in nature, is solved with aid of SDSC supercomputers and published in *Science*.

KidLab, an after-school program for 10-12 year olds, established at SDSC.

CERFnet officially dedicated at SDSC; Vinton Cerf (no relation to the network) conducts dedication.

CRAY Y-MP8/864 supercomputer arrives at SDSC and made available December 22.

1989

Charles Keeling, et. al., with SDSC supercomputers, construct 3-D model of the terrestrial carbon cycle confirming the importance of fossil fuel combustion in loading the atmosphere with carbon dioxide.

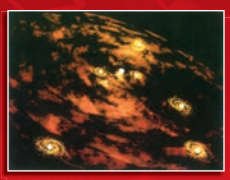

Dave Nadeau and Holliday Horton at SDSC create an animation of accretion disks – the fiery nebulae and spinning clouds in a primordial solar system – for a planetarium show at S.D.'s Reuben H. Fleet Space Center and Theatre.

1990

SDSC acquires 32-node iPSC/860 parallel computer from DARPA, built by Intel.

SDSC acquires 256-node NCUBE 2 parallel computer.

1991

UCSD/SDSC Researchers Solve Structure for the Body's "Transistors"

Protein kinases have been likened to cellular regulatory circuits in living organisms that perform similar functions as transistors or chips in a computer. In 1991, a team of researchers from SDSC and UCSD reported in a cover story in *Science* that they had solved its three-dimensional structure; the solution—with the aid of the CRAY supercomputer at SDSC and a stereoscopic visualization system in SDSC's Advanced Scientific Visualization Lab (VizLab)—was considered one of the Grand Challenges of computational and biological science. Researchers around the globe then began searching for specific kinase inhibitors that target specific diseases, including diabetes and tumor initiation and growth.

1992

SDSC Takes Big "STEP" with Innovative Program for Local Teachers

With the beginning of the 1992 school year, SDSC launched its Supercomputer Teacher Enhancement Program (STEP) and presented its first half-day-in-service to 25 educators and parents from Grant Math Science Magnet School. The program—a forerunner to SDSC's award-winning TeacherTech program—was designed to show teachers how scientists use computers to make discoveries, and it introduced teachers to ways elementary school students could use computers to learn about math and science.

1995

Catching an Alleged Cyber-criminal

On February 15, 1995, SDSC Senior Fellow Tsutomu Shimomura and Systems Analyst Andrew Gross collaborated with federal agents to track down alleged computer criminal Kevin Mitnick. Shimomura and Gross applied their knowledge of computer network security to help agents apprehend Mitnick, then considered the "most-wanted computer criminal in the United States," after an intruder broke into a network of computer systems at Shimomura's home and at SDSC. Mitnick, now a computer consultant and author, was convicted of various computer- and communications-related crimes; he was released from prison in January 2002.

UCSD/SDSC researchers solve the structure for protein kinase, likened to the body's transistor, and considered one of the Grand Challenges in biological sciences; research makes cover of *Science*.

SDSC becomes first site to send messages cross country through the NSFNET T-3 backbone; at 45 Mb/s, it's the fastest openly available network for research and education.

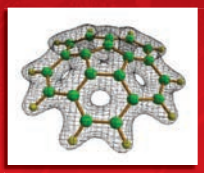

Computer studies of corannulene, a bowl-shaped molecule that constitutes one-third of the molecular shell of the Bucky Ball, is subject of cover story of *Chemical and Engineering News*; research by Kim Baldridge and Jay Siegel from SDSC/UCSD.

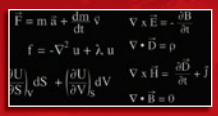

SDSC unveils "Grand Challenge Equations" exhibit.

Intel Paragon supercomputer arrives at SDSC.

1991

UC receives a 3-year, $15 million grant from DEC to develop an advanced information and data management system to increase the productivity of researchers studying global change – called Project Sequoia.

1992

First workshop to discuss issues related to the National Metacenter, a synthesis of the intellectual and computational resources of the four NSF supercomputering centers — SDSC, Cornell Theory Center, NCSA, and PSC — was held at SDSC.

SDSC launches STEP (Supercomputer Teacher Enhancement Program) to develop education outreach programs for K-12 with local educators.

1993

CRAY C90 supercomputer arrives at SDSC; officially installed during November press briefing.

TOP 30

1998

Catching a Speeding Enzyme in the Act

With large-scale computer simulations run at SDSC, researchers showed how one of the fastest enzymes–acetylcholinesterase (AChE), which controls communications among nerves and muscle cells–works. The speed of AChE had been puzzling, since its active site appeared to be accessible only by a partly blocked channel on the enzyme's surface. Earlier work showed that "breathing" motions in AChE open and close the channel to allow acetylcholine (ACh) to enter the active site. The new work showed that the breathing motions allowed ACh to bind almost as fast as if the channel were always open. A A team of researchers, including J. Andrew McCammon (UC San Diego) combined computational models and theoretical calculations to obtain their results, published in the August 4 *Proceedings of the National Academy of Sciences.*

1998

SDSC Releases Glue that Holds Data Together

In 1998, the Storage Research Broker (SRB) 1.1 was released as the "middleware" that holds together data cache sites for NPACI, the National Partnership for Advanced Computational Infrastructure. The SRB software, built on the work of Reagan Moore at SDSC, is still used by many U.S. and international computational science research projects. It is considered a "middleware" in the sense that it is built on top of other major software packages (various storage systems, real-time data sources, etc.) and it has callable library functions that can be used by higher-level software. SDSC's Chaitan Baru, Michael Wan, Arcot Rajasker and Wayne Schroeder were members of the original team that developed SRB.

1999

Computer Simulations Reveal New Anti-HIV Strategy, Leading to AIDS Drug

Molecular dynamics simulations, conducted by a team led by UCSD chemist J. Andrew McCammon, provided new insights into attacking a third target against HIV—integrase—that helps the virus hijack the body's cells. The simulations, published in 1999 in the *Biophysical Journal*, led to the development of Isentress, marketed by Merck as a new HIV drug approved for patient use by the FDA. Hailed as the most important new AIDS drug in a decade, the drug was the first AIDS medicine to block integrase, considered crucial in the process HIV uses to replicate.

A data transfer speed record of 630 Mb/s is achieved across the 100-mile CASA Gigabit Testbed link between SDSC and Caltech, accelerating solution of the reaction of atomic hydrogen with molecular heavy hydrogen (deuterium) by a factor of 3.3.

Chris Mihos and Lars Hernquist of UC Santa Cruz collaborate with computer artists at NCSA, using computational resources at SDSC CRAY C90, to create high-resolution images of a galaxy encounter for IMAX cosmic voyage, which debuts at the Smithsonian National Air and Space Museum in D.C.

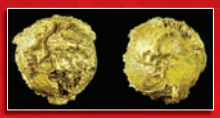

SDSC establishes the telemanufacturing facility to rapidly prototype 3D models from digital geometry data.

CRAY T3E supercomputer installed at SDSC.

SDSC receives $8.4M contract from DARPA to develop Distributed Object Computation Testbed (DOCT) for handling complex documents on geographically distributed data archives and computing platforms; to focus on the needs of the US Patent and Trademark Office.

1994

NIH approves $3.286M to SDSC to fund the National Biomedical Computation Resource (NBCR).

Thinking Machines CM-2 arrives at SDSC to support UCSD education and research.

1995

For the first time, SDSC harnesses the power of a new very high-speed network (vBNS) by distributing portions of a computation across high-performance computers located on the east and west coasts.

Tsutomu Shimomura, SDSC Senior Fellow, collaborates with federal agents to track down alleged "notorious cybercriminal" Kevin Mitnick, then considered the "most-wanted computer criminal in the United States."

1996

A model of the nicotinic acetylcholine receptor is developed by Igor Tsigelny, Naoya Sugiyama and Palmer Taylor at UCSD/SDSC, in collaboration with Steven Sine at the Mayo Foundation; enzyme is a target for addictive activity from nicotine.

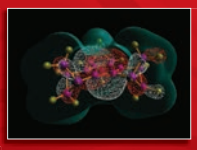

The cover of *Chemical and Engineering News* features an image of cyclohexatriene molecule; research results from computational and experimental collaboration of Kim Baldridge at SDSC and Jay Siegel, at UCSD.

1999

World's Largest Repository of Protein Structures Housed at UCSD/SDSC

On July 1, 1999 responsibility for the Protein Data Bank (PDB)–the world's largest archive for biomedical structures–formally shifted to the Research Collaboratory for Structural Bioinformatics (RCSB) with a new PDB website and ftp archive. Today, the program is managed jointly by two partner sites: Rutgers University, under the direction of Helen Berman; and SDSC and the Skaggs School of Pharmacy and Pharmaceutical Sciences at UCSD. The PDB houses about 100,000 biological structures, including proteins associated with the common cold, avian flu, HIV, West Nile virus, Alzheimer's disease, and a wide variety of cancers.

2000

Transporting Theater-goers at Hayden Planetarium to the Orion Nebula

Astronomically accurate visualizations made possible for the first time by SDSC researchers and the IBM *Blue Horizon* supercomputer at SDSC transported space theater visitors to the Orion Nebula—the first destination of the virtual starship departing from the reopened Hayden Planetarium at the American Museum of Natural History in New York. The visualizations were made possible by SDSC's Galactic MPIRE volume-rendering software package, under the technical leadership of SDSC's Dave Nadeau.

2000

SDSC Technologies Provide Huge Image Archive to Study Human Embryology

SDSC-developed storage and visualization technologies were integrated into a National Library of Medicine project to create one of the largest-ever medical image databases. The project, called "Human Embryology Digital Library and Collaboratory Support Tools," was designed to demonstrate how leading-edge information technologies in computation, visualization, collaboration, and networking can expand the capabilities of medical science in developmental studies, clinical work, and teaching. The database allowed project participants to study data sets of sizes up to a terabyte, in multi-gigabyte images, using the IBM HPSS archival storage system at SDSC, the Storage Resource Broker, and the MPIRE 3-D system to support the 3-D rendering of data.

SDSC named leading-edge site for the National Partnership for Advanced Computational Infrastructure (NPACI), launched in October.

CAIDA established with a NSF seed grant to promote a more robost, scalable Internet infrastructure; principal investigator, kc Claffy.

The world's largest and most powerful transmission electron microscope is operated from UC San Diego and the National Center for Microscopy and Image Research (NCMIR) in a successful demonstration of trans-Pacific telemicroscopy by American and Japanese researchers.

The Research Collaboratory for Structural Bioinformatics (RCSB), under the management of Helen Berman at Rutgers University, and Phil Bourne at UCDS/SDSC, assumes primary responsibility for the Protein Data Bank — the world's largest archive for biomedical structures used in pharmacological and medical research.

1997

SDSC publishes "Women in Science," featuring bios of women who had a career in, or made significant contributions, to a scientific discipline.

1998

The Storage Research Broker (SRB) 1.1 is released as "middleware" to hold together data cache sites from NPACI, led by SDSC. The SRB software is built on work led by SDSC's Reagan Moore; Chaitan Baru, Michael Wan, Arcot Rajasekar and Wayne Schroeder are members of the original team that developed SRB.

With large-scale computer simulations run at SDSC, researchers led by J. Andrew McCammon at UCSD show how one of the fastest enzymes in the world, acetylcholinesterase, does its work; results are published in the *Proceedings of the National Academy of Sciences*.

1999

The HPC Systems group completes the first installation of a user file system that is greater than one terabyte in capacity.

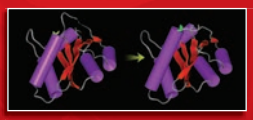

Molecular dynamics simulations, led by J. Andrew McCammon at UCSD, show how to attack a third enzyme target integrase, a crucial enzyme used by HIV to replicate; research led to the development of Isentress, an anti-AIDS drug marketed by Merck.

TOP 30

2001

On the "Rocks"– Cluster Computing Made Simple

Under the leadership of Phillip Papadopoulos at SDSC, the National Partnership for Advanced Computational Infrastructure (NPACI) designed and released an enhanced version of the NPACI Rocks Cluster Toolkit, a set of open-source enhancements for managing Linux-based computer clusters. Over time, Rocks—simple, self-contained, scalable and upgradable—slowly became the *de facto* software package for implementing clusters. By the end of 2001, NPACI Rocks was used to establish clusters at the Pacific Northwest National Laboratory, Northwestern University, the University of Texas, and Caltech.

2001

TeacherTech Launched, Bringing Computer Science to the Classroom

SDSC organized the first San Diego TeacherTech during the summer of 2001, a program designed to help educators bring new technology tools and technology-enabled science concepts into K-12 curricula. Within five years, the program attracted more than 1,200 teachers from more than 150 area schools to its workshops; it was estimated that these educators reached as many as 200,000 students from San Diego County and Baja, Mexico. In 2006, SDSC TeacherTECH was presented with a Partner of the Year Award by the San Diego Science Alliance.

2002

SDSC Experts Help Reveal Impact of Climate Change on Mexico

In a study published in *Nature*, a team of researchers from SDSC, the University of Kansas, and the Universidad Nacional Autonoma de Mexico (UNAM) reported they had analyzed the potential impacts of climate change on the ecosystems of Mexico–the first such analysis for entire country–including 1,870 species of mammals, birds, and butterflies. The interdisciplinary study was made possible with the help of a powerful software program, the Genetic Algorithm for Rule-Set Prediction (GARP), created by SDSC's David Stockwell.

Animated visualizations made possible for the first time by SDSC's Galactic MPIRE volume-rendering software and *Blue Horizon* are displayed in the all new Hayden Planetarium of the American Natural History Museum on New Year's Eve.

IBM's *Blue Horizon* is delivered to SDSC as the most powerful computer available to the US academic community – capable of 1 trillion FLOPS.

Using sophisticated "backscatter analysis", CAIDA researchers track the progress of a worm dubbed Code-Red Worm, which infected hundreds of thousands of Web servers around the world.

SDSC organzies the first San Diego TEACHERTECH, designed to allow San Diego K-12 teachers to explore and examine different multimedia and web applications for use in the classroom.

"The Search for Life: Are We Alone?" – a new space show at the Hayden Planetarium in New York – premiers to rave reviews. SDSC plays a key role in creating a realistic animation showing the birth of our solar system

2000

A map showing possible paths an email message might take in the Internet, created by the skitter tool developed at CAIDA at SDSC, makes the cover of *Nature*.

SDSC-developed storage and visualization technologies – integrated into a National Library of Medicine project – create one of the largest-ever medical image databases.

NSF awards $2.3 million, three-year grant to UCSD to create, demonstrate, and evaluate a non-commercial prototype, high-performance wide-area wireless network for research and education (HPWREN)—a multi-institutional collaboration led by Hans-Werner Braun at SDSC and Frank Vernon at SIO.

2001

NSF establishes TeraGrid to support world-class scientific discovery and education through a grid-based cyberinfrastructure; SDSC is one of the founding sites.

Fran Berman becomes director of SDSC.

A team led by Vincent Crespi at Penn State used computer simulations and resources at SDSC to discover carbon fibers with mechanical strength comparable to a diamond – strong and stiff carbon tubes called nanotubes.

2002

The first analysis of the potential impacts of climate change for an entire country, Mexico, is reported in a paper published in *Nature* by a team that included researches from SDSC.

2004

In the Beginning...

Michael Norman, professor of physics at the Center for Astrophysics and Space Sciences (CASS) at UC San Diego, together with colleagues at CASS and SDSC, ran the world's largest and most complex scientific simulation of the evolution of the universe ever performed. Using SDSC's IBM Blue Horizon supercomputer, the team tracked the formation of enormous structures of galaxies and gas clouds during the millions and billions of years following the Big Bang. Norman ran his "Enzo" cosmology program for more than 100 hours on all 128 computing nodes of the *Blue Horizon*.

2004

Envisioning the "Big One" for Southern California

A collaboration of 33 earthquake scientists, computer scientists, and others from eight institutions produced the largest and most detailed simulation yet of just what might happen during a major earthquake–magnitude 7.7–on the southern San Andreas Fault. The simulation, known as TeraShake, used the new 10 teraflops IBM *DataStar* supercomputer and large-scale data resources of SDSC. The simulation provided more detail into how intensely the earth would shake during such an event, and what impact it would have on structures, particularly in the populated sediment-filled basins of Southern California and northern Mexico.

2006

Predicting "Solar Storms"

At times, "solar storms" ejected from the sun's corona–it's ghostly outer atmosphere–can eject plasma in the direction of the Earth, resulting in potentially serious disruptions in satellite operations, communications, and even electrical power grids. Since society is heavily dependent on these infrastructures, predicting "solar storms" are of tremendous importance. The March 29, 2006 solar eclipse gave scientists from the Solar Physics Group at SAIC (Science Applications International Corporation) an opportunity to check their predictions of the state of the solar corona based on a computational model using observed photospheric magnetic field data. Using dedicated time on SDSC's IBM supercomputer *DataStar* and NASA's *Columbia* system, this work represented the most true-to-life computer simulation ever made of the solar corona.

PRAGMA (Pacific Rim Applications and Grid Middleware Assembly), launched during a workshop at SDSC and funded by NSF, shows how relationships and expertise developed to tackle computational research could also help thousands of SARS patients in Taiwan.

Mike Norman, and colleagues at Center for Astrophysics and Space Sciences at UCSD, run the world's largest and most complex scientific simulation of the evolution of the universe ever performed.

SDSC's High Performance Storage System (HPSS) reaches the milestone of one petabyte of stored data.

Data experts at SDSC collaborate with American Red Cross to help locate missing loved ones in the wake of Hurricane Katrina; results in "Safe and Well" website.

CENIC announces that the first production 10 Gigabit Ethernet campus connection in the U.S. has been installed from UCSD.

2003

The TeraGrid enters production with two clusters installed at SDSC: IBM/Intel IA-64 TeraGrid Phase 1 Cluster and IBM/Intel IA-64 TeraGrid Phase 2 Cluster.

2004

Scientists led by UCSD's J. Andrew McCammon use molecular simulations and SDSC resources to identify a potential mechanism underlying the drug resistance of the worst mutant HIV strain; in same work, the researchers identify a separate region of protease enzyme that might serve as new drug target.

SDSC is the first academic institution in the world to install the new IBM eServer Blue Gene Solution computing system.

2005

SDSC receives $2.2M award from NIH to provide Next Generation Biology Workbench, building on the work of the "Workshop" concept developed by Shankar Subramaniam at UCSD/SDSC.

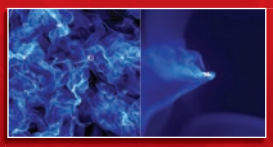

Astrophysicist Richard Klein from UC Berkeley and others use simulations run at SDSC to explode one of two competing theories about how stars form inside immense clouds of interstellar gas; results published in *Nature*.

SDSC launches DataCentral, the first program of its kind to support large community data collections and databases.

TOP 30

2006

Predicting Protein Structure in Record Times

Researchers from SDSC, contributing their massive computational capabilities to a collaboration with colleagues at the University of Washington and IBM, helped achieve the largest-ever protein structure prediction–and completed the complex simulation in less than three hours, a task that previously took weeks. The ground-breaking demonstration used UW Professor David Baker's Rosetta Code and ran on more than 40,000 central processing units of IBM's *Blue Gene Watson* supercomputer, using the experience gained on the IBM's *Blue Gene* system at SDSC. Ross Walker, a SAC computational scientist at SDSC, managed the Baker's group access to SDSC machines, helping them to optimize their code.

2006

HPWREN Comes to the Aid of Local Fire Fighters

Firefighters facing fast-spreading wildfires, especially in remote areas where communications and other resources are scarce, added "cyberinfrastructure" to their fightfighting arsenals during the 2006 "Horse Fire" in California's Cleveland National Forest. Experts from the High Performance Wireless Research Educational Network (HPWREN)—a resource supported by the NSF and staffed by researchers at SDSC, Scripps Institution of Oceanography, and San Diego State—responded to the urgent request of state firefighters for quick and reliable wireless communications among widespread teams. Within a day, HPWREN experts were on the scene, establishing high-speed wireless data links needed to contain the flames.

2007

Zeroing in on the Cause of Alzheimer's and Parkinson's Disease

Using the massive computer-simulation power of SDSC, researchers at UC San Diego homed in on the causes of Alzheimer's disease, Parkinson's disease, and other neurological disorders. A cover story in the March issue of the *Federation of European Biochemical Societies Journal (FEBS)* offered, for the first time, a model for the complex process of aggregation of a protein known as alpha-synuclein, which in turn, leads to harmful ring-like or pore-like structures in human membranes—the kind of damage found in Parkinson's and Alzheimer's patients. The researchers also found that the destructive properties of alpha-synuclein can be blocked by beta-synuclein—a finding that could lead to treatments for many debilitating diseases.

NARA and SDSC, with concurrence from NSF, sign a landmark MOU that provides an avenue for preserving valuable digital data collections; first time NARA establishes an affiliated relationship for preserving digital data with an academic institution.

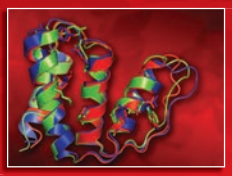

Researchers at SDSC — working with colleagues at the University of Washington — achieve the largest-ever protein structure prediction and complete the simulation in less than three hours.

UCSD/SDSC researchers zero in on causes of Parkinson's disease, Alzheimer's disease, and other neurological disorders with a computer model featured on cover of the *Federation of European Biochemical Societies Journal (FEBS)*.

SDSC releases version 0.5 of iRODS, the open-source Integrated Rule-Oriented Data System, which represents a new approach to digital data management.

The source of spider silk's strength, as strong as steel, is simulated by MIT scientists in collaboration with applications scientist Ross Walker at SDSC, on SDSC's IBM Blue Gene/L supercomputer.

A team led by Laura Carrington at SDSC successfully completes a record-setting, petascale-level simulation of the earth's inner structure; a finalist for Gordon Bell Prize.

2006

Firefighters facing fast-spreading wildfires urgently request cyberinfrastructure resources from HPWREN to help combat the "Horse Fire" in Cleveland National Forest.

The most true-to-life computer simulation ever made of our sun's corona — created by researchers at Science Applications International Corp., with the help of SDSC resources — successfully predicted its actual appearance during the total solar eclipse of March 29.

2007

A team of researchers from NCAR, SDSC, LLNL and IBM Watson, led by Allan Snavely at SDSC, set U.S. records for size, performance, and fidelity of computer weather simulations, modeling the kind of "virtual weather" that society depends on for accurate weather forecasts; a finalist for Gordon Bell Prize.

Science is coming to the YouTube generation with the advent of "SciVee," a collaboration between the NSF and SDSC, under the direction of Phil Bourne, UCSD/SDSC.

2008

SDSC dedicates a new, energy-efficient building extension as a key resource for UC San Diego and beyond.

CAIDA researchers Dmitri Krioukov and kc Claffy, along with Marián Boguñá (Universitat de Barcelona), reveal in *Nature Physics* a previously unknown mathematical model called "hidden metric space" that may explain the "small world phenomenon," offering a potentially more efficient way to pass messages on the Internet.

2007
Setting Records for "Virtual Weather" Prediction

A team of researchers from SDSC, the National Center for Atmospheric Research (NCAR), Lawrence Livermore National Laboratory, and the IBM Watson Research Center set U.S. records for size, performance, and fidelity of computer weather simulations, modeling the kind of "virtual weather" that society depends on for accurate weather forecasts. The research, led by Allan Snavely at SDSC, was a finalist for the Gordon Bell Prize.

2008
Removing Bottlenecks from the Internet

SDSC/CAIDA researchers Dmitri Krioukov and Kimberly Claffy, along with Marián Boguñá (Universitat de Barcelona), reveal in *Nature Physics* a previously unknown mathematical model called "hidden metric space" that may explain the "small world phenomenon" — discovered in the 1960s by sociologist Stanley Milgram—and its relationship to natural and manmade networks. The novel concept suggests a new approach for removing bottlenecks within the Internet that threaten the smooth passage of digital information around the globe.

2008
Modeling Earth's Enigmatic Core

To learn more about the inner sanctum of the earth's core, seismologists take advantage of one of nature's most destructive forces: earthquakes. Somewhat like the way a CAT scan images the brain, seismologists track seismic wave patterns from earthquakes to model the structure of the earth's core. One of the great challenges is to capture the propagation of high-frequency waves, with periods of 1 to 2 seconds, as they travel across the globe. In 2008, a team of researchers from SDSC employed a spectral-element application called SPECFEM3D_GLOBE to complete record-setting, petascale-level simulations of the earth's core at frequencies over just over a second. Results of these simulations were announced at SC08, where the research led by Laura Carrington of SDSC was a finalist for the Gordon Bell Prize.

SDSC awarded 5-year, $20 million grant from NSF to build *Gordon*, a powerful supercomputer featuring "flash memory" and "supernodes" to solve critical data-intensive science problems.

SDSC unveils *Dash*, a "flash-memory-based" supercomputer to accelerate solutions for data-intensive science problems.

The CIPRES portal, used to help researchers track evolutionary relations among species, becomes the most heavily used portal in the TeraGrid, accounting for 20% of active TeraGrid users during the first quarter of 2010.

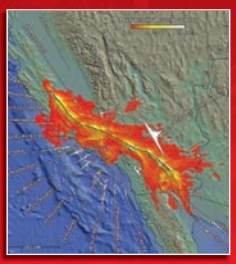

Researchers at SDSC, SDSU and UCSD create the largest-ever simulation of a Magnitude 8 earthquake, primarily along the southern section of the San Andreas fault.

Trestles, an HPC system designed to offer modest-scale and gateway users rapid job turnaround to increase researcher productivity, is launched under a $2.8 million NSF grant.

The Center for Large-scale Data Systems Research (CLDS), bringing together industry and university research to investigate "big data" challenges, is launched under the direction of SDSC researcher Chaitan Baru.

2009

SDSC officially launches the *Triton Resource*, an integrated data-intensive computing system primarily designed to support UCSD and UC.

SDSC completes a comprehensive upgrade to its tape-based archival storage capacity, increasing its total to 36 petabytes, the largest digital storage capacity of any academic center in the world.

2010

SDSC launches a volunteer internship program for high school students – Research Experience for High School Students (REHS) – to help them gain experience in a particular area of computational research.

Mike Norman named SDSC director.

2011

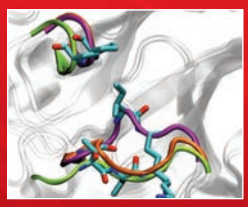

Researchers from UC Irvine, led by Rommie Amaro and using SDSC expertise and resources, find a new approach to create customized therapies for virulent flu strains that resist current antiviral drugs. The findings were published in *Nature Communications*.

TOP 30

2009

Providing a Portal to Build the Tree of Life

To help scientists build the Tree of Life—to infer the evolutionary history of Earth's myriad species starting from biomolecular sequence data—the NSF funded a project called CIPRES (CyberInfrastructure for Phylogenetic RESearch). As part of this project, SDSC developed the CIPRES portal—a browser interface to the most widely used phylogenetic codes—along with faster versions of these codes. In December 2009 the portal was migrated to the TeraGrid, the nation's largest open-access network of high-performance computers. Within the first quarter of 2010, the portal had 500 users, the most of any TeraGrid portal and 20 percent of all active TeraGrid users. Each month since then, more than 100 new users have accessed TeraGrid resources through the portal, attesting to the broad impact of this enabling interface.

2011

Using Simulations to Create Customized Therapies for Virulent Flu Strains

The search for effective flu drugs has always been hampered by the influenza virus itself, which mutates from strain to strain, making it difficult to target with a specific pharmaceutical approach. Researchers from UC Irvine, with assistance of SDSC expertise and computer resources, found a new approach to create customized therapies for virulent flu strains that resist current antiviral drugs. The findings, published in *Nature Communications*, offered an avenue to build new drugs that exploit so-called flu protein 'pockets.' Using powerful computer simulations on SDSC's *Trestles* system, UCI's Rommie Amaro and Robin Bush with SDSC's Ross Walker created a method to predict how pocket structures on the surface of influenza proteins promoting viral replication could be identified as these proteins evolve, allowing for possible pharmaceutical exploitation.

2012

Internet Censorship Revealed Through a Maze of Malware

In January 2011, Egypt—with 23 million Internet users—vanished from cyberspace after its government ordered an Internet blackout amidst anti-government protests. The following month, the Libyan government, also under siege, imposed an Internet "curfew" before completely cutting access for almost four days. To help explain how these governments disrupted the Internet, a team of scientists led by kc Claffy, director and founder of CAIDA at SDSC, conducted an analysis based largely on the drop in a specific subset of observable Internet traffic that is a residual product of malware. Their analysis, funded by the NSF and Department of Homeland Security—and including scientists from Italy and The Netherlands—was the first published research to demonstrate how malware-generated traffic pollution could be used to analyze Internet censorship and other macroscopic network outages.

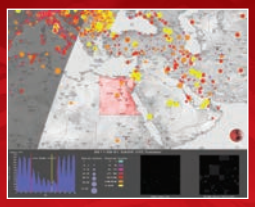

To help explain how governments disrupt the Internet, a team of scientists led kc Claffy, director and founder of CAIDA at SDSC, conducts the first published research showing how malware-generated traffic pollution can be used to analyze Internet censorship and other network outages.

A study led by SDSC's James Short predicts that by 2015, the sum of media delivered to consumers on mobile devices and to their homes would take 15+ hours a day to consume. That's equal to nine DVDs worth of data per person per day.

The NSF awards SDSC a $12 million grant to deploy *Comet*, a new petascale supercomputer designed to transform advanced scientific computing by expanding access and capacity among traditional as well as non-traditional research domains.

Researchers from SDSC, the U.S. Geological Survey, and the San Diego Zoo's Institute for Conservation Research develop methodology that for the first time combines 3D and advanced range estimator technologies, providing detailed data on the movements of terrestrial, aquatic, and avian wildlife species.

2012

The Predictive Analytics Center of Excellence (PACE), is funded by SDSC to leverage the Center's data-intensive expertise and resources to provide the critical skills to design, build, verify, and test predictive data models.

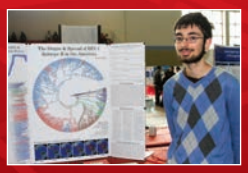

Arman Bilge, a 10th grader at the Lexington High School in Massachusetts, used the CIPRES science gateway, developed by SDSC researchers, to create a map and timeline that identified when HIV arrived in Americas, and where and when it spread across North and South America.

2013

WIFIRE, designed to build a cyberinfrastructure to perform real-time data-driven assessments of wildfires, is funded under a three-year NSF grant to SDSC, Calit2/Qualcomm Institute, and the Jacobs School of Engineering's MAE department. Also participating is the University of Maryland's Department of Fire Protection Engineering.

2014

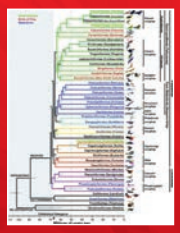

A published global genome study using SDSC's data-intensive *Gordon* supercomputer have researchers rethinking how avian lineages diverged after the extinction of the dinosaurs. The four-year project, called Avian Genome Consortium, is published in the journal *Science* in late 2014.

2013

U.S. Media Intake to Rise to 15.5 Hours per Day, per Person, by 2015

A study led by a researcher at SDSC predicted that by 2015, the sum of media asked for and delivered to consumers on mobile devices and to their homes would take more than 15 hours a day to see or hear. That volume is equal to 6.9 million-million gigabytes of information, or a daily consumption of nine DVDs worth of data per person per day. "One can actually have more than 24 hours in a media day," explained James E. Short, the author of the latest "How Much Information?" report, produced by the Institute for Communications Technology Management (CTM) at the University of Southern California's Marshall School of Business. "As we increase our level of multi-tasking, we have to expect that total hours will grow even as the total number of physical hours a viewer can consume media will remain roughly constant."

2014

Genomic Analysis of 115 Year-old Woman

Researchers investigating the genome of a healthy supercentenarian discovered many somatic mutations that arose during the woman's lifetime. However, very few mutations mapped to regions in genes that code for proteins, whereas the overwhelming majority were in genomic regions predicted to have neither adverse nor favorable impact on genetic fitness. Led by Erik Sistermans and Henne Holstege from the VU University Medical Center in Amsterdam, the team—including SDSC researchers Wayne Pfeiffer and Mark Miller—published its findings in the journal *Genome Research*. The work involved numerous computations, some of which were done on the *Triton* cluster at SDSC. At the time of her death at the age of 115, the subject woman was the second oldest person in the world and showed no signs of vascular disease or dementia.

2014

Rethinking How Our Feathered Friends Evolved

A published global genome study that used SDSC's data-intensive *Gordon* supercomputer had researchers rethinking how avian lineages diverged after the extinction of the dinosaurs. The four-year project, called Avian Genome Consortium and published in the journal *Science* in late 2014, resulted in a new family "tree" for nearly all of the 10,000 species of birds alive today by comparing the entire DNA codes (genomes) of 48 species as varied as parrot, penguin, downy woodpecker, and Anna's hummingbird. The massive undertaking, started in 2011, involved more than 200 researchers at 80 institutions in 20 countries, with related studies involving scientists at more than 140 institutions worldwide.

About this Brochure

On November 14, 1985, the San Diego Supercomputer Center opened its doors on the northwestern corner of the UC San Diego campus and showed off its first supercomputer, a CRAY X-MP/48 — clocking at what one newspaper article called a "mind-boggling" billion calculations per second... a gigaflop. Actually, the peak performance was closer to 800 million calculations per second, but still pretty fast then. Some 100 researchers — all from traditional disciplines such as astrophysics, biochemistry, geology, and oceanography — applied for time on the new supercomputer, which promised to usher in a new era of scientific discovery.

Since that day, scientific and technological advances made possible, and/or created by SDSC staff and resources like the original CRAY, have made a major mark in academia, industry, and society-at-large – "turning data to discovery," a phrase that has become associated with SDSC. The Center has brought together researchers at UC San Diego and across the nation and world, in partnerships and collaborations that now are the hallmark of today's scientific enterprise. SDSC also has proven to be a good neighbor, providing its expertise and considerable resources to local educators and students, firefighters and other "first-responders," families of military serving overseas, and others in time of need.

To help commemorate SDSC's 30th anniversary, SDSC has pulled together a timeline of the most significant events in the Center's history. This timeline, now located outside SDSC's data center in the building's East Wing addition, serves as a physical reminder of the historical milestones that have made SDSC a local, state, and national resource for high-performance and data-intensive computing, and a leader in research and development of the nation's vast cyberinfrastructure.

As an added element to this celebration, this brochure spotlights 30 significant moments and/or advances in science, technology, and outreach made possible by the Center and its staff. The "Top 30" list, published in this document, represents but a small sampling of the hundreds of major accomplishments over SDSC's history; clearly, many other achievements also deserve recognition. We are confident this list will be discussed and debated, and ultimately revised and updated when the next such list is developed!

SDSC

San Diego Supercomputer Center
University of California, San Diego
www.sdsc.edu | twitter/SDSC_UCSD
facebook/SanDiegoSupercomputerCenter