UNIVERSITY OF SOUTHERN CALIFORNIA

EXCELLENCE ACROSS THE DISCIPLINES

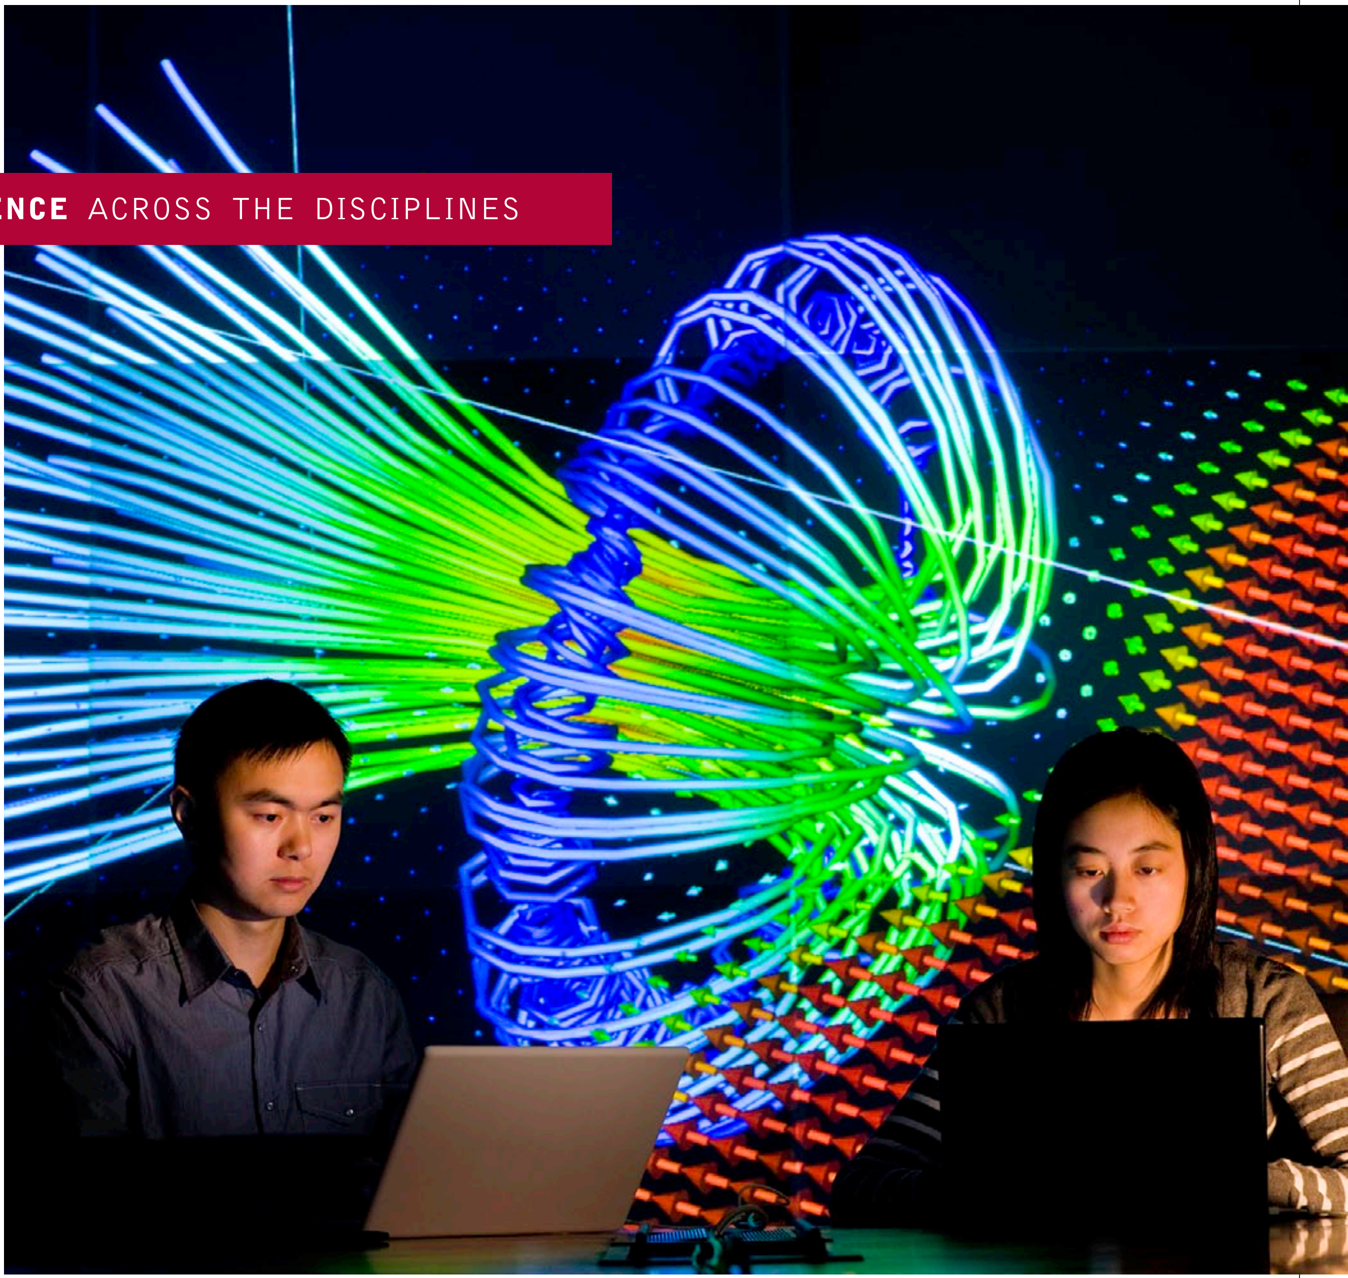

CENTER FOR HIGH-PERFORMANCE COMPUTING AND COMMUNICATIONS **(HPCC)** **2011**

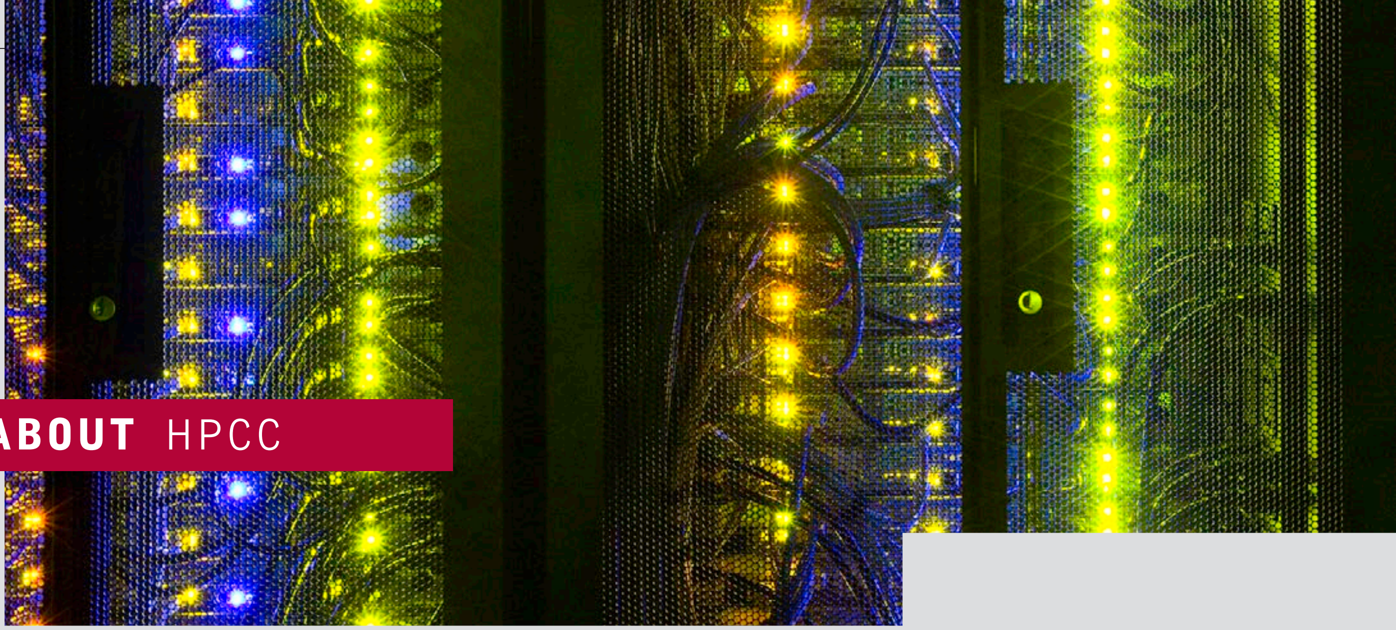

ABOUT HPCC

The Center for High-Performance Computing and Communications (HPCC) serves as a foundation and model for computational research at the University of Southern California. HPCC bridges USC's unique strengths in scientific computing, computer science, and communications by supporting research groups in a variety of disciplines, including biophotonics, computational biology, cosmology, geophysics, linguistics, materials science, and pharmacology.

HPCC resources enable researchers to move beyond the limitations of traditional research. A prime example of this is the work of Michelle Povinelli, assistant professor in the USC Viterbi School of Engineering's Ming Hsieh Department of Electrical Engineering and holder of the Women in Science and Engineering Junior Gabilan Chair. Povinelli and her colleague Andrea Armani, featured on page 6, both had the distinction of receiving a 2010 Presidential Early Career Award for Scientists and Engineers, the highest honor bestowed by the United States government on outstanding young scientists and engineers.

Povinelli's research explores nano- and micro-photonic device technology, which has applications in optical fiber communications, solar energy, and biomedicine. Povinelli and her research team study the optical properties of nanostructured materials, such as photonic crystals, metamaterials, and micro-resonators, which are patterned on the submicron scale—100 times smaller than the diameter of a human hair. Using the same technology that is used to make computer chips, Povinelli's team fabricates the devices by etching tiny holes or lines in silicon.

Povinelli develops devices to process information in optical fiber communication networks with the aim of improving data transmission over the Internet. Povinelli and her team also study how nanostructured materials can be used to boost light absorption in next-generation solar cells. In their recent work, they investigate the use of light forces to probe cell biomechanics and the self-assembly of complex materials. HPCC resources enable Povinelli to perform detailed simulations of the flow of light through nanostructured materials. The large-scale, highly parallel architecture of HPCC resources is essential for performing her team's computationally intensive tasks of simulating realistic 3-D structures and optimizing device designs.

The overarching goal of HPCC is to expand the reach of "big computing" by addressing system-level questions in social, economic, and cultural research. With its advanced computing and communications resources, HPCC helps to create the virtual organizations, virtual worlds, and immersive environments that will continue to transform and globalize higher education in the twenty-first century.

Elizabeth Garrett
Provost and Senior Vice President for Academic Affairs

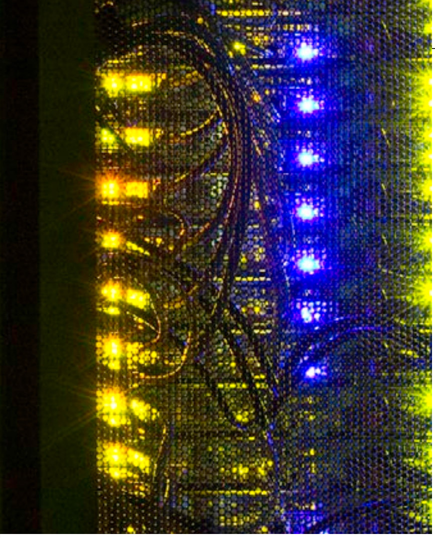

Launched in 2000, the University of Southern California's Center for High-Performance Computing and Communications (HPCC) is a global leader in research computing.

Among supercomputers in an academic setting, USC's HPCC supercomputer cluster is the 5th fastest in the United States and the 17th fastest in the world. Among all supercomputers in the world, USC's HPCC cluster is ranked 63rd. It claimed this distinction by achieving a benchmark in spring 2011 of 126.40 teraflops, or 126.40 trillion floating-point calculations per second, on its 1,988-node, quad-core and hex-core, dual-processor cluster on a 10-gigabit Myrinet backbone.

USC's HPCC has achieved its current status among the world's top supercomputer sites through local investments and without national funding. The local aspect of HPCC allows USC faculty and graduate students unfettered access to a world-class facility, enabling them to tackle even the most challenging data-intensive scientific problems. HPCC helps graduate students develop the twenty-first century skills necessary to become global leaders in advanced technologies and critical scientific discoveries.

Priya Vashishta
Professor of Chemical Engineering and Materials Science,
Computer Science, and Physics and Astronomy
Faculty Executive Director, USC Center for
High-Performance Computing and Communications

COVER: Weiwei Mou and Ying Li, graduate students in materials science in USC's Mork Family Department of Chemical Engineering and Materials Science, study the velocity streamlines of a nanojet formed after a shock wave has induced the collapse of a water nanobubble. **PAGE 02:** Rich Seymour, graduate student in materials science in USC's Mork Family Department of Chemical Engineering and Materials Science, examines the progression of slip planes and dislocations in a simulation of nanoindentations on a nickel-aluminide surface. **PAGE 03:** Amy Zaoshi Yuan, a graduate student in USC's Mork Family Department of Chemical Engineering and Materials Science, visualizes siRNA, a double stranded RNA molecule noted for its role in RNA interference.

USC CENTER FOR DATA **VISUALIZATION** AND **COLLABORATION**

Located in the Seaver Science Center, the Center for Data Visualization and Collaboration houses technologies that provide researchers with the means to teleconference, collaborate, and share complex information in high-resolution visual forms.

Powerful tools for visualization give USC faculty and graduate students the ability to view HPCC-generated simulation data from different perspectives and perform large-scale, in-depth analyses requiring specialized equipment in a relatively risk-free setting.

The center includes a tiled display wall and an AccessGrid that provides video, audio, and networking equipment with high-speed access to distributed data and computers. This equipment may also be linked to support group-to-group interactions across USC's network and the Internet. The room contains three 65-inch, movable LCD screens that are capable of displaying one continuous image or three separate images.

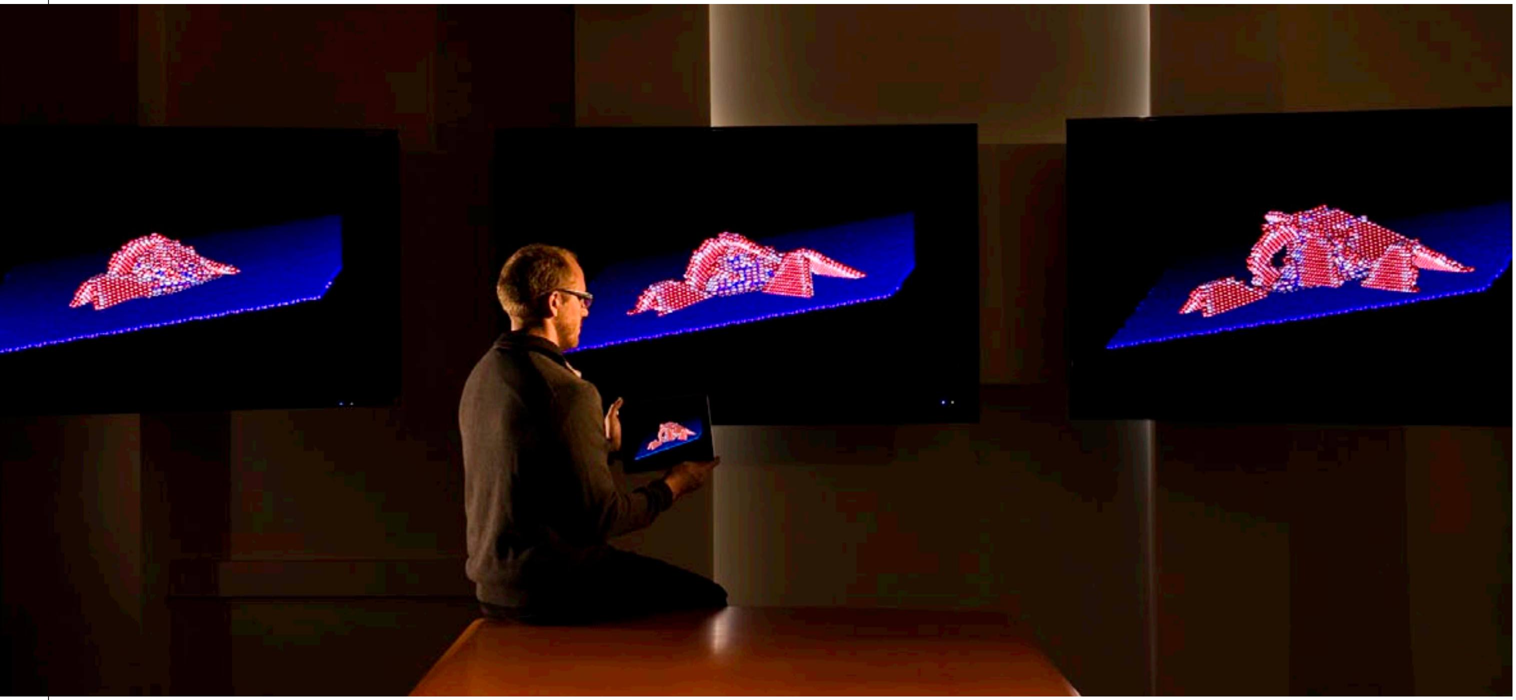

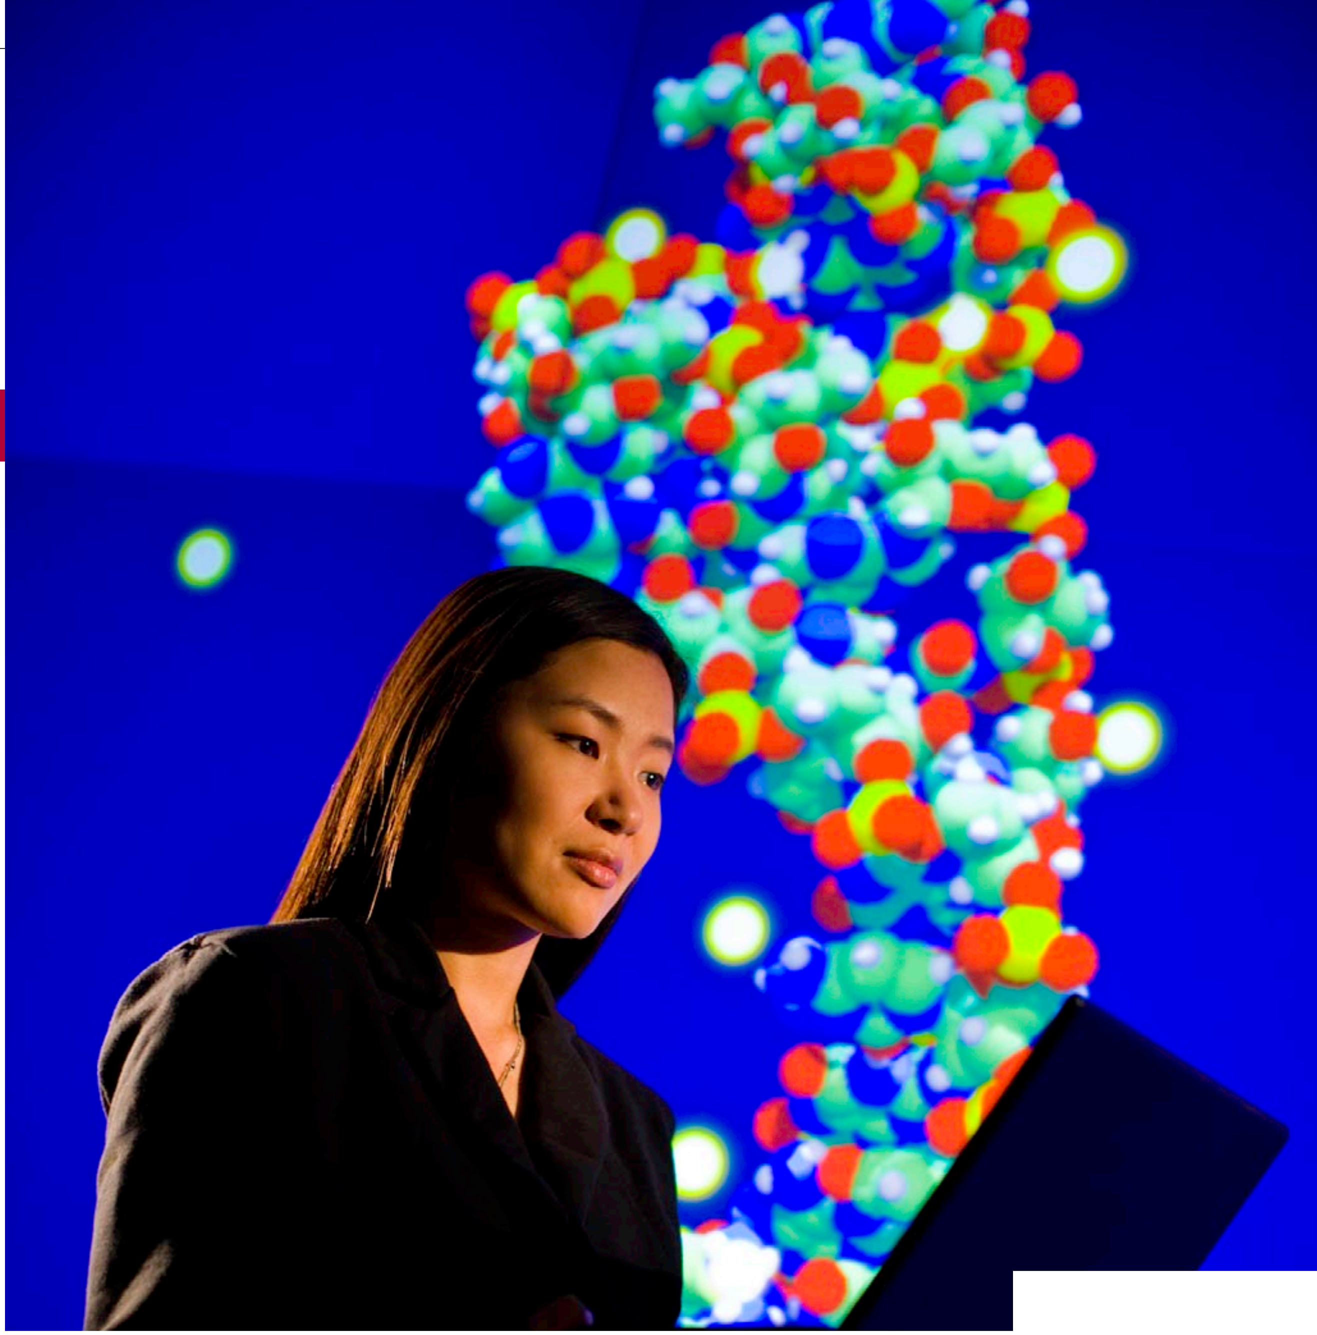

AT&T Tile Wall

The AT&T tile wall provides users with a scalable means of displaying wall-sized images or computer simulations with very high resolution. Moving or still digital images are sliced up into a dozen smaller images by a small cluster of computers and then projected by a dozen digital projectors onto a 14-foot-wide and 8-foot-high, monolithic, all-glass screen. The rear-projected composite image, which contains roughly 17.5 million pixels, has remarkable clarity when viewed close up or at a distance.

AccessGrid

The AccessGrid allows users to videoconference and collaborate interactively over the Internet. Three projectors, suspended from the ceiling, display computer simulations, films, live television feeds, or websites onto a wall-sized screen. Participants from a variety of locations—ranging from university classrooms and labs to hotel conference rooms and government offices—can come together virtually for meetings, collaborative work sessions, seminars, lectures, and training.

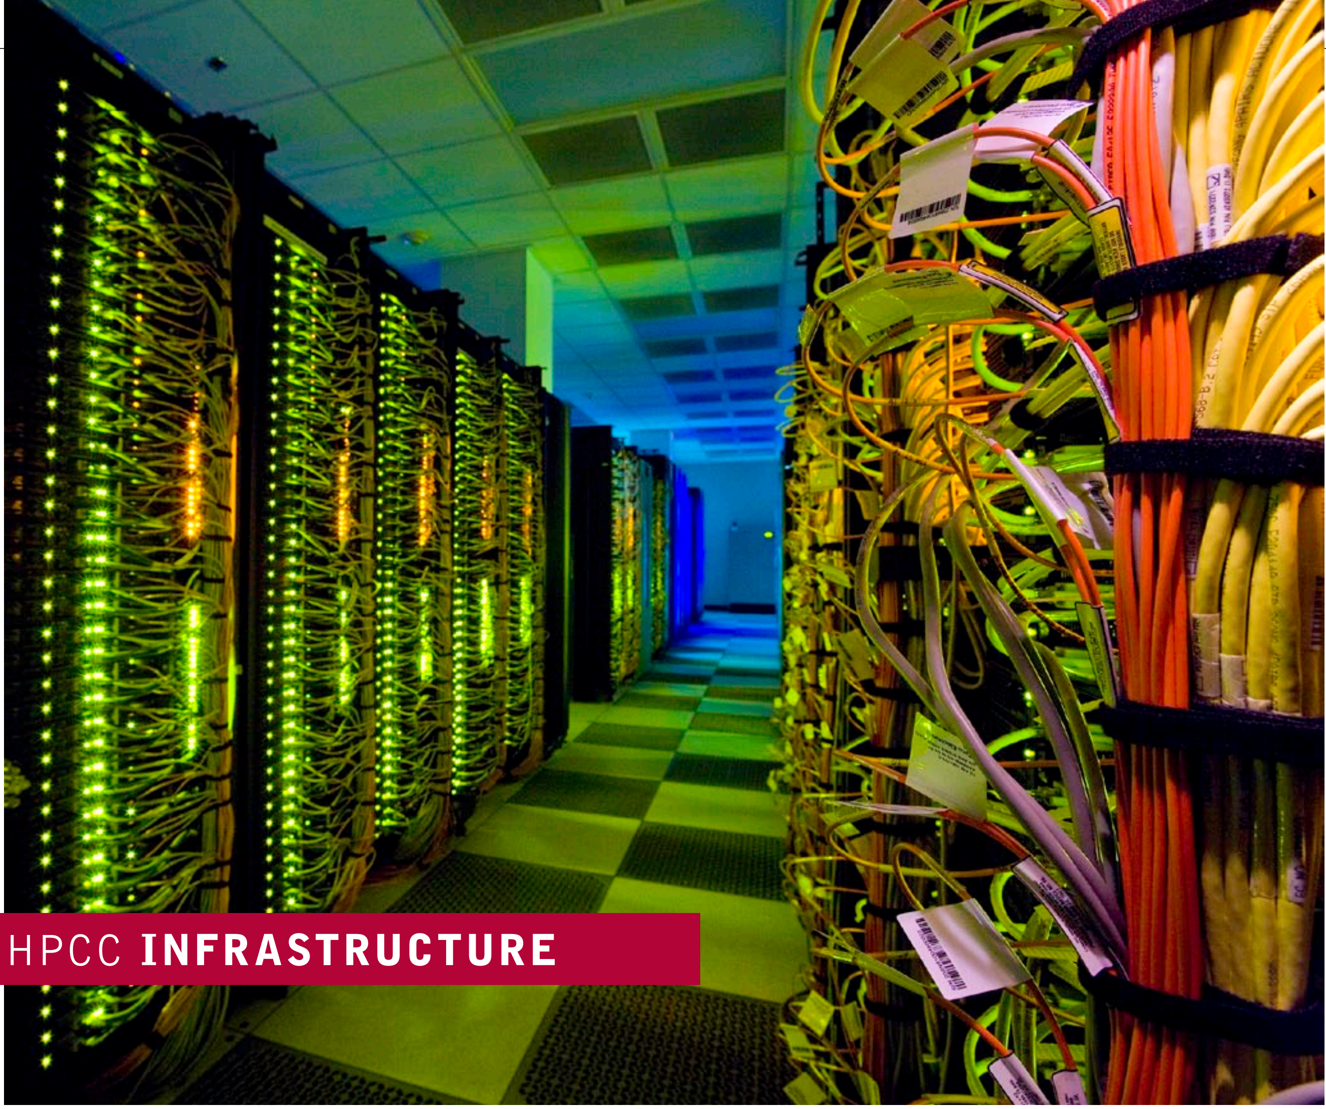

HPCC **INFRASTRUCTURE**

HPCC comprises a diverse mix of computing and data resources. Two Linux clusters constitute the principal computing resource. In addition, HPCC has a central facility that provides more than 550 terabytes of combined disk storage.

Housed in USC's state-of-the-art data center, HPCC allows faculty members to maximize their individual research resources by contributing to HPCC's condominium-style compute environment. Under this arrangement, researchers may use their grant awards to acquire HPCC nodes, which are accessed through the same interface as the standard HPCC nodes. HPCC staff members monitor and maintain the equipment on a round-the-clock basis, enabling researchers to focus on their primary research. Researchers have sole access to their condo nodes except for twice a year when HPCC runs the LINPACK benchmarks.

Linux Clusters

HPCC has two Linux clusters. The 1,988-node, quad-core and hex-core, dual-processor cluster contains Dell, Oracle Sun, Hewlett-Packard, and IBM nodes on a 10-gigabit Myrinet backbone. The second cluster contains 787 dual-processor nodes—of which nearly half are dual-core AMD Opteron™ processors—on a 2-gigabit Myrinet network. The combined main memory of the two Linux clusters is more than 31 terabytes; their combined temporary disk storage is more than 490 terabytes. The compute power of HPCC enables USC researchers to tackle grand-challenge class problems in a broad range of scientific disciplines.

For each cluster, the bidirectional, low-latency Myrinet fiber network interconnects the nodes, allowing for the development of massive production jobs that require high-speed communications among computational elements.

The 2-gigabit Linux cluster also provides 5 large-memory nodes with 64 gigabytes of RAM and 8 dual-core AMD Opterons.

Networks

HPCC has multiple-gigabit interfaces to the USC campus network. These interfaces facilitate massive data transfers and interactive access to the computing facilities at 10 gigabits per second, creating a campus grid environment through USC's multiple-gigabit-per-second campus network.

HPCC is home to the headquarters of Los Nettos, the regional network that provides redundant and reliable network bandwidth to the Los Nettos Consortium, comprising USC, the California Institute of Technology (Caltech), the Jet Propulsion Laboratory, Loyola Marymount University, and the Claremont Colleges.

Now in its 23rd year of operation, Los Nettos' infrastructure of dark fiber and leased circuits uses leading-edge optical technologies to enhance the network's flexibility and provide services such as private virtual local area networks (VLANs), dynamic interconnects, and partitioned wavelengths between member sites.

Through its affiliations with the Corporation for Education Network Initiatives in California (CENIC), Internet2, National LambdaRail, and Internet exchanges such as Pacific Wave, CIIX, and Any2, Los Nettos offers high-capacity access to many national and international research and education networks.

Pacific Wave

Pacific Wave has developed a distributed Internet exchange facility on the West Coast (currently in Seattle, Sunnyvale, and Los Angeles) to allow interconnection among U.S. and international research and education networks. The Pacific Wave exchange, operated by Pacific Northwest Gigapop (PNWGP) and the Corporation for Education Network Initiatives in California (CENIC), enables data to pass directly between major national and international networks, increasing the efficiency and speed of data transfer while eliminating the costs associated with routing data through multiple circuits and dedicated interconnects. This infrastructure supports a wide variety of advanced science and engineering applications. USC researchers have access to this Internet exchange and its international and national participants through the Los Nettos regional network.

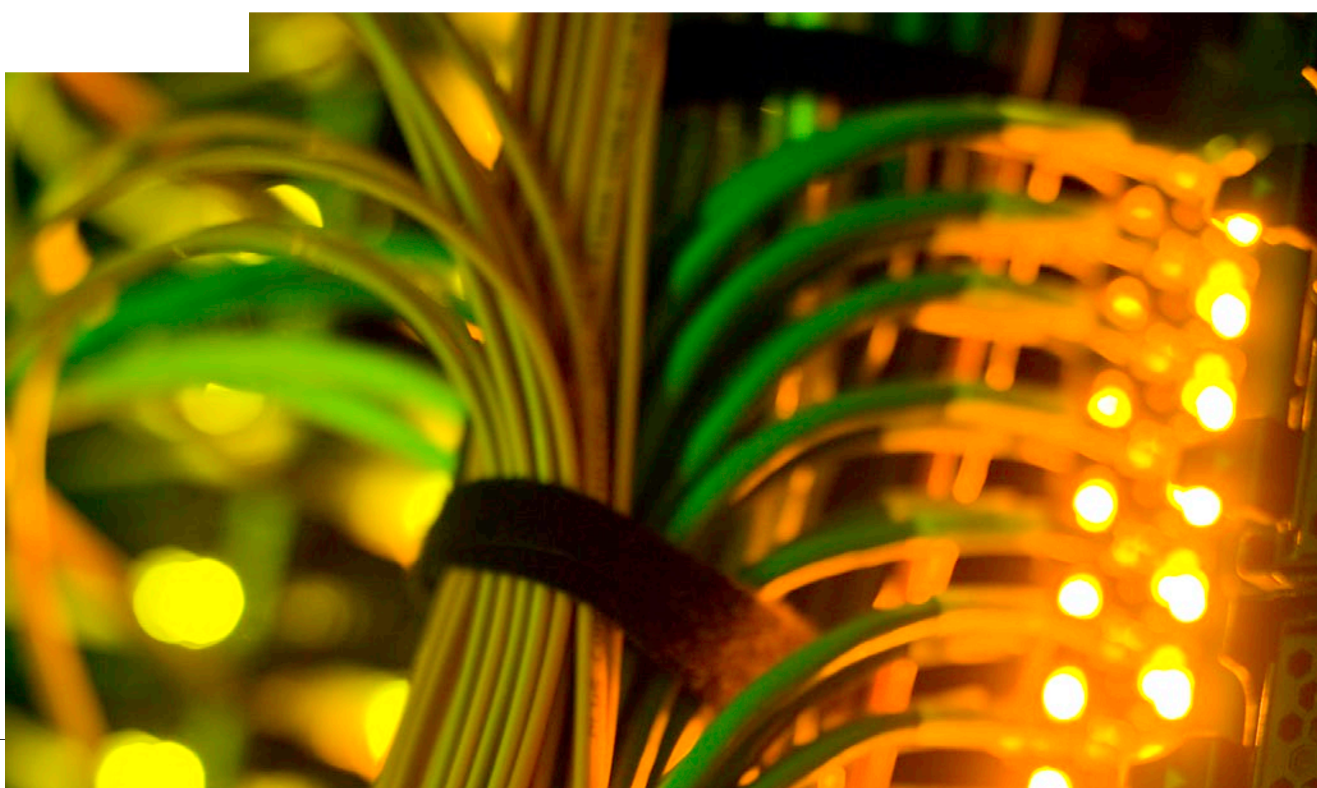

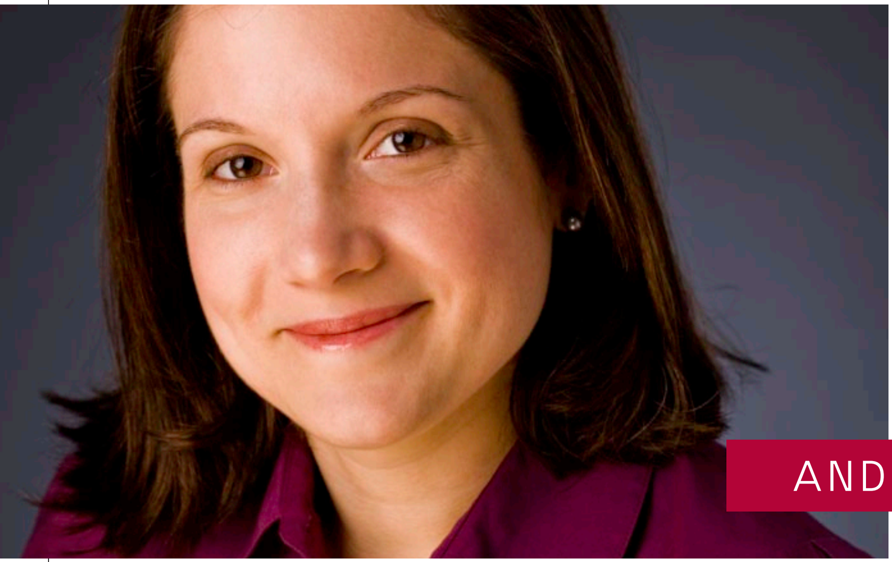

ANDREA **ARMANI**

Andrea Armani is an assistant professor in the USC Viterbi School of Engineering's Mork Family Department of Chemical Engineering and Materials Science and holds the Fluor Early Career Chair of Engineering. Her research explores biophotonics and nanotechnology, from integrated photonics to optical biosensors, and enhances the understanding of the physics of biological and chemical materials.

Armani and her research group develop new optical materials and devices, such as waveguides and lasers, which can be employed as sensing platforms. New optical materials under development by Armani's group include silica glasses doped with rare-earth metals and quantum dots. Her lab is also creating surface modification techniques to improve how optical sensors work.

Armani is recognized for demonstrating the first optical sensor capable of detecting single molecules without relying on labels, an achievement with important applications in medicine. In contrast to conventional fluorescent diagnostic methods, which detect at discrete time intervals, like a photograph, label-free methods continuously monitor and provide information, like a video-feed. This is a significant advantage, both in disease diagnosis and treatment monitoring.

As the next step in the field, Armani and her group are designing a nanolaser-sensor to identify the triggers that cause cancer. The device will detect the changes in DNA as they happen in real time, making it possible, for the first time, to study a single DNA strand in isolation, rather than having to focus on hundreds or thousands of strands. She and her group hope their device will help identify which chemical triggers people should avoid in order to prevent cancer.

Armani's research is funded by the Department of Defense, National Institutes of Health, and National Science Foundation. She has received numerous awards, including the Office of Naval Research Young Investigator Award, the Technology Review Top 35 Innovators under 35 Award, the USC Mellon Mentoring Award for Undergraduate Mentoring, the National Institutes of Health New Innovator Award, and the Presidential Early Career Award for Scientists and Engineers.

BELOW: A rendering of a silica waveguide splitter integrated on a silicon wafer. This optical device, which transfers light from a single input port into a secondary waveguide, has applications ranging from optical computers to portable biosensors.

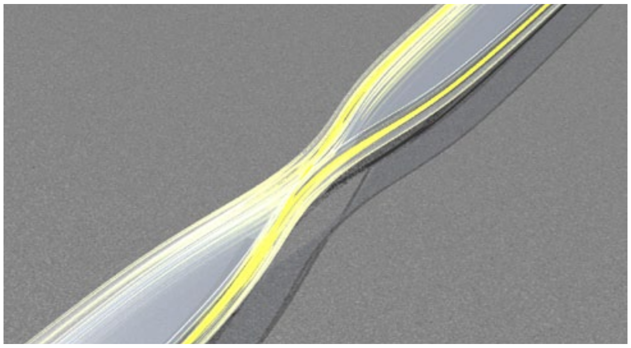

STEPHEN **CRONIN**

Stephen Cronin is an assistant professor in the USC Viterbi School of Engineering's Ming Hsieh Department of Electrical Engineering and holds the Gordon S. Marshall Early Career Chair in Engineering. He and his research team study photocatalytic water splitting, a process that has applications for renewable energy, including the production of hydrogen fuel. His team also explores the photocatalytic production of methane and other hydrocarbon fuels from carbon dioxide and water.

Photocatalytic processes, which employ light to drive chemical reactions, provide a method for storing the sun's energy in chemical bonds that can be released later without producing harmful byproducts. Cronin's team is developing new methods and materials to enhance these photocatalytic processes.

When absorbed by a semiconductor, light creates electron-hole pairs that are separated in energy by the bandgap of the material. This energy separation can be used to drive electrons in a circuit, as in an electrical solar cell, or to drive electrochemical redox reactions, such as water splitting or the formation of methane from carbon dioxide and water. Cronin and his team are combining strongly plasmonic metal nanoparticles, such as gold, with strongly catalytic metal oxide semiconductors, such as titanium dioxide, to enable more efficient solar energy conversion.

Cronin's team uses HPCC resources to perform detailed simulations of the electromagnetic response of plasmonic metal nanoparticles on semiconductor photocatalysts. These simulations provide a model for studying the fundamental mechanism underlying the observed catalytic enhancement, as well as the necessary data to calculate and visualize electromagnetic fields and charges in these metal and semiconductor nanostructures. The large-scale, highly parallel architecture of HPCC's resources is essential for simulating these structures, which typically require a spatial resolution of two angstroms and several million grid points.

Cronin's research is funded by a National Science Foundation Faculty Early Center Development (CAREER) Award and an Air Force Office of Scientific Research Young Investigator Award.

BELOW: This figure shows the electromagnetic response of a gold nano-island film deposited on a titanium-dioxide photocatalyst. The red spots show localized regions of intense electromagnetic activity due to the plasmonic response of the film.

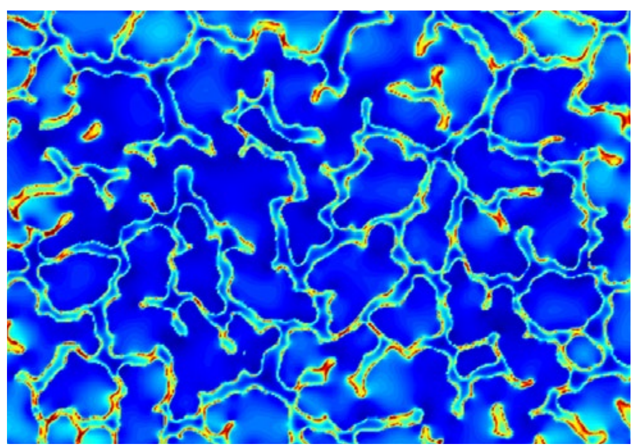

NOURI **NEAMATI**

Nouri Neamati is an associate professor in the USC School of Pharmacy's Department of Pharmacology and Pharmaceutical Sciences. His research interests include drug design and discovery, database mining of small-molecule compounds, molecular pharmacology, and chemoinformatics.

Neamati has established an internationally recognized program in drug design and discovery with a major focus on antiviral and anticancer therapeutics. He and his research team use computational chemistry to match chemical compounds to biological processes in order to design drugs for the treatment of various diseases, including cancer and HIV/AIDS. Using the structural information available for many existing drugs and targets, Neamati and his team develop small-molecule compounds based on known drugs and the size and shape of a targeted cavity within a protein of a diseased cell. The compound must conform closely to the shape of the cavity in order to bind to it and inhibit the function of the protein.

Neamati's team uses HPCC resources to test thousands of compounds against the shapes of single, diseased cells to determine the best fit. His team has developed a database, containing 10 million small-molecule compounds, that is readily searchable using two-dimensional fingerprinting algorithms. The team has also calculated up to 200 conformations for each compound to generate approximately 2 billion structures that are fully searchable in three dimensions.

Neamati and his team are currently testing numerous compounds that could lead to major breakthroughs in the treatment of cancer and HIV/AIDS.

Neamati's research is funded by the National Institutes of Health (NIH), Department of Defense, American Lung Association, American Association for Cancer Research, and the Susan G. Komen Breast Cancer Foundation. He has been awarded the NIH Technology Transfer Award, the Stop Cancer Award, the GlaxoSmithKline Drug Discovery and Development Award, and an Idea Award from the Department of Defense's Congressionally Directed Ovarian and Breast Cancer Research Program.

BELOW: A simulation of a small-molecule compound (pink) binding to a targeted cavity site (blue) within a glucose-regulated GRP-78 protein (purple).

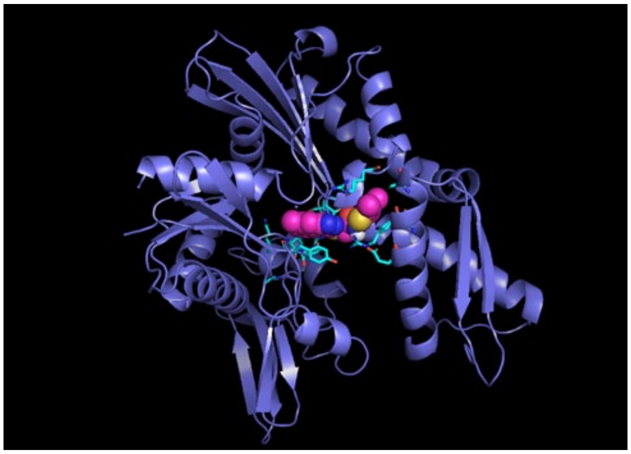

SIMON **TAVARÉ**

Simon Tavaré holds the George and Louise Kawamoto Chair in Biological Sciences and is a research professor in Molecular and Computational Biology in the USC Dornsife College of Letters, Arts, and Sciences. Tavaré applies the power of mathematical sciences to help solve complex biological and medical problems through collaborations with paleontologists, pathologists, cancer genomicists, statistical geneticists, molecular biologists, and population geneticists.

One of the pioneers in the field of computational biology, Tavaré has made key contributions in statistical bioinformatics and the development of evolutionary genomic approaches for understanding cancer. Tavaré's best-known work centers on the use of DNA sequence data to trace the lineage of a cell, an individual, or a species back through time.

Tavaré's group uses HPCC resources to study aspects of tumor evolution. He and his research team use variation in the DNA of tumor cells to infer a tumor's ancestral history. To generate the molecular data, they have developed a high-throughput, experimental technology for assessing variability in millions of individual tumor cells. Approximate Bayesian computation methods are then used to infer parameters of biological interest. One goal of this research is to identify the role that cancer stem cells play in the evolution of tumors.

Tavaré is also the principal investigator of the USC Center of Excellence in Genomic Science (CEGS). Under Tavaré's leadership, the center is investigating the process by which genotypic variation translates into phenotypic variation. The CEGS uses HPCC resources for the analysis of next-generation sequencing data, including DNA, RNA, and methylation data. The center hopes to create a unified picture of how different genetic variants interact with the environment to influence aspects of disease.

Tavaré's research is funded by grants from the National Institutes of Health, the National Science Foundation, and Cancer Research UK. Tavaré is a Fellow of the United Kingdom's Academy of Medical Sciences and the Royal Society.

BELOW: A 3-D simulation of a cellular Potts model of a tumor. Each region, indicated by a color, represents a type of mutation present in the tumor.

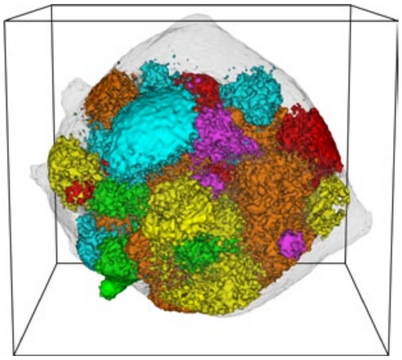

UNIVERSITY OF SOUTHERN CALIFORNIA

CENTER FOR HIGH-PERFORMANCE COMPUTING AND COMMUNICATIONS **(HPCC)**

www.usc.edu/hpcc

PRINCIPAL **INVESTIGATORS** AND **PROJECTS** SUPPORTED BY HPCC

Hooman Allayee	Genome-Wide Association Study of Complex Diseases
Murali Annavaram	Computer Systems Reliability Using WearMon
Andrea Armani	Modeling of Optical Devices
Prithviraj Banerjee	Human Motion Detection and Tracking for Video
James Baurley	Algorithm for Pathway Structures
Thorsten Becker	Global Mantle Convection and Crustal Dynamics
Theodore Berger	Large-Scale Hippocampal Model
Todd Brun	Quantum Error Correction for Orbital Angular Momentum in Photons
Nikolaus Buenning	Investigating the Predictability of Drought
Jacqueline Chame	Parallel Programming
Gary Chen	Simulation of Genetic Variation in Admixed Populations
Keith Chin	Atomistic Simulation of Protein/Surface Interaction and Protein/Surfactant Interaction
David Conti	Statistical Methods Development in Genetic Epidemiology
Stephen Cronin	Numerical Simulation of Nanoscale Systems
Enoch Dames	Combustion Kinetics Modeling of Jet Fuel Surrogates
Pedro Diniz	Auto-Tuning DSL
Julian Domaradzki	Numerical Simulations of Turbulent Flows
Fokion Egolfopoulos	Aerodynamic and Kinetics Processes in Homogeneous and Heterogeneous Reacting Flows
Julien Emile-Geay	Climate Change Data and Detection: Modeling the Tropical Climate of the Past Millennium
William Gauderman	Genome-Wide Association Study of Childhood Respiratory Outcomes
Panayiotis Georgiou	Speech-to-Speech Translation
Roger Ghanem	Algorithms and Applications for Uncertainty Quantification
Aman Goel	Semantic Annotation Using Conditional Random Fields
Norberto Grzywacz	Visual Processing in the Brain
Sam Gustman	Conversion of the USC Shoah Foundation Institute Video Testimonies into Multiple Digital Formats
Stephan Haas	Magnetic Fields and Disorder in Quantum Spin Systems
Samuel Hartzmark	Behavioral Finance
John Heidemann	Los Angeles Network Data Exchange and Repository (LANDER) Project
Thieo Hogen-Esch	Elucidation of the Mechanisms of the Yamamoto Polymerizations
Eduard Hovy	Semantics-Oriented Natural Language Processing Research
Manish Jain	Stackelberg Equilibria
Dinakar Jayarajan	Creating Dynamic Models of Promotions
Sewoong Jung	Particle Tracking Using the Lattice Boltzmann Method
Shina Kamerlin	Computer Simulations of Chemical and Biological Systems
Samuel Kim	Analysis of Unstructured Audio Data Including Environmental and Physiological Signals
Kevin Knight	Statistical-Based Natural Language Processing
James Knowles	Analysis of DNA Sequences Generated with Next-Generation Technology
Dayung Koh	Mathieu Equation
Georgios Konstantinidis	Query Reformulation for Health Data Integration
Anna Krylov	Electronic Structure Calculations
C.C. Jay Kuo	Multimedia Technologies
Manish Kurse	Structure and Function of the Fingers' Tendinous Apparatus
Jun-Young Kwak	Execution-Time Reasoning in Distributed Partially Observable Markov Decision Processes
Darius Lakdawalla	Insurance Benefit Design
Jiin-Jen Lee	Computational Fluid Dynamic Analysis of Bridges Exposed to Hurricane Waves
Mihye Lee	Analysis of European Firm-Level Data
Natasha Lepore	Morphometry in MRI images of the Brains of Infants and Fetuses
Kristina Lerman	Social Networks
Yiyu Li	Molecular Modeling of Protein Structures
Yan Liu	Multi-Aspect Anomaly Detection
Zhenshuo Liu	Simultaneous CO_2 Recycling in an *in situ* Combustion Operation
Thomas Ly	Case-Control Simulations
Philip Maechling	Cyberinfrastructure for Earthquake System Science
Donald Marsh	Nephron Synchronization
Patric Muggli	Full-Scale Modeling of High-Energy Particle Accelerators
Aiichiro Nakano	Ultrascale Atomistic Simulations of Chemical Reactions
Shrikanth Narayanan	Tracking Dynamics of Vocal Tract Shaping
Nouri Neamati	Virtual Screening of Large Databases of Small-Molecule Drugs
Michael Neely	Optimizing Coordinated Combination Drug Therapy
Feng Ni	Binding Free Energy for Drug Design of New Farnesyl Pyrophosphate Synthase Inhibitors
Ken-ichi Nomura	Multimillion Atom Simulations of Energetic Materials
John O'Brien	Large Finite-Difference Time-Domain Simulations of the Electromagnetic Fields in Microcavity Light Emitters
Nicos Petasis	Molecular Modeling for Drug Discovery
Elena Pierpaoli	Cosmological Parameter Determination
Krzysztof Pilch	Phases of N=8 Supergravity
Michelle Povinelli	Large-Scale Simulations of Microphotonic Devices
Viktor Prasanna	Semantic Web Technologies for Smart Oil Fields
Peter Qin	Molecular Dynamics Analyses of Nitroxides Attached to Nucleic Acids
Lars Rahm	Synthesis of Novel High-Nitrogen Molecules
Joaquin Rapela	Responses of Visual Cells to Naturalistic Stimuli
Golam Rasul	Ab Initio/GIAO-MP2/CCSD(T) Calculations of Structures and NMR and Other Properties of Highly Reactive Carbocations, Carbodications, Onium Ions, and Dications
Edward Rhodes Jr.	Helioseismic Studies of Solar Internal Structure and Dynamics
Robert Sacker	Modeling the Formation of E. Coli Morphotypes
Muhammad Sahimi	Landfill Modeling and Estimation
Takahiro Sakai	Numerical Simulation of Hydrodynamic Model in a Stratified Lake
Abhishek Sharma	Job Scheduling for Data Centers
Junyang Shen	Development of Wireless Technology for Location Estimation
Darko Stefanovski	Steered Molecular Dynamics of P-type ATPases
Alexander Stram	Increasing Accessibility to 1000 Genomes' Data
Daniel Stram	Genome-Wide Association Scans
Brian Sutch	Development of a Protein-Nucleic Acid Interaction Database for Potential Use in Drug Design Against Any Protein Target
Rei Tangko	Combustion Chemistry
Alexander Tartakovsky	Multidisciplinary University Research Initiative Air Force Project
Simon Tavaré	Computational Biophysics and Biochemistry
Paul Thomas	Methods for Reconstructing Gene Evolution
Barry Thompson	Stereoregular Electroactive Polymers
Priya Vashishta	Multimillion-to-Billion Atom Molecular Dynamics Simulations of Materials
Tom Vernier	Molecular Dynamics Simulations of Biomolecular Structures in Nanosecond, Megavolt-per-Meter, Pulsed Electric Fields
Greg Ver Steeg	Methods for Detecting Clusters
Fang Wang	Asymmetric Catalysis
Hao Wang	Cross-Linguistic Speech Analysis in Children
Joseph Wang	Particle Simulation
Kai Wang	Next-Generation Sequencing Data Analysis
Arieh Warshel	Computational Biophysics and Biochemistry
Richard Watanabe	The BetaGene Study: Identifying Genes Predisposed for Insulin Resistance
Joanne Wu	Patterns of Care and Costs for the Treatment of Rheumatoid Arthritis in a Medicaid Population
Shanshan Xu	Multivariate Regression Analysis
Maocai Yan	Molecular Modeling of the APOBEC3G Protein
Liye Zhang	Computer Simulation of Kondo Effect
Xiaoxin Zhang	Analyzing the Impact of Mortgage Agencies on Residential Mortgages with Competing Risk Framework

DESIGN: WARREN GROUP | STUDIO DELUXE PRINCIPAL PHOTOGRAPHY: JOHN LIVZEY