Information Sciences Institute

USCViterbi
School of Engineering

USC Viterbi
School of Engineering

Information Sciences Institute

Information Sciences Institute (ISI)

is a world leader in research and development of advanced information processing, computing and communications technologies.

Creativity is thinking up new things.
Innovation is doing new things."
— Theodore Levitt

During its 45 years of operation, ISI researchers have enabled a number of fundamental advances in the domain of information and computing sciences. Examples include the design and implementation of the Internet's domain name system (DNS), voice-over-IP technology, fabless foundries, polymorphic robots, statistical machine translation, rhetorical structure theory, trusted electronics, characterizing the behaviors of individuals and communities in social media, grid computing, experimental cybersecurity research, and adiabatic quantum computing. As I begin my fourth year as the executive director of ISI, I feel privileged to lead an organization which has an incredibly impactful history of scientific and technological contributions, and holds equally exciting prospects for the future.

This report is a celebration of the accomplishments and vitality of the ISI community. As you peruse it, you will become acquainted with ISI's research staff, their research and academic activities, professional community contributions, and recognitions they have received. Through the descriptions of research carried out in our centers and divisions, our collaborations with industry, and specific scientific highlights, you will gain insight into the dynamism, diversity, and impact of ISI's research — and its continued focus on investigating and solving present-day, real-world critical issues such as combating the scourge of human trafficking through the development and deployment of innovative new technologies. In summary, you will find that ISI embodies a unique synthesis of creativity and innovation.

One measure of an organization's vitality is its financial health. In 2015 our annual budget was $73.3M and growing. This included funding from government agencies such as DARPA, IARPA, DHS, ONR, NIH, NSF, ARO, and DOE; partnerships with academic institutions such as Carnegie Mellon, Rutgers, Cornell, Stanford, Caltech, and MIT; and partnerships with industrial organizations such as Lockheed Martin, Northrop Grumman, Raytheon, Leidos, MITRE, Next Century Corporation, and Charles River Analytics.

Another measure of ISI's vibrancy, and arguably the more important one, is the talent that it attracts. In 2015 we continued to attract top researchers and students in the fields of computer science and engineering. Well over one hundred doctoral, master's and undergraduate students receive invaluable research experience at ISI every year. And this past year, twelve of our graduate students in scientific areas that include network security, ocean tomography, data integration, situational awareness/intrusion detection, and natural language processing received their doctorates and accepted positions at major research institutions such JPL, Sandia National Labs, USC, and Oxford University, as well as leading technology companies such as Google, Microsoft, NVIDA, Facebook, and Yelp.

In closing, on behalf of ISI, I wish to extend our appreciation and gratitude to Michael and Linda Keston for their generous philanthropic gift in the fall of 2015. This gift allowed us to establish the Keston Executive Directorship Endowment, which will be used to further our program of excellence in research, and specifically, to recruit and fund new, talented researchers-in-residence.

Premkumar Natarajan
Michael Keston Executive Director, ISI/USC
Vice Dean of Engineering
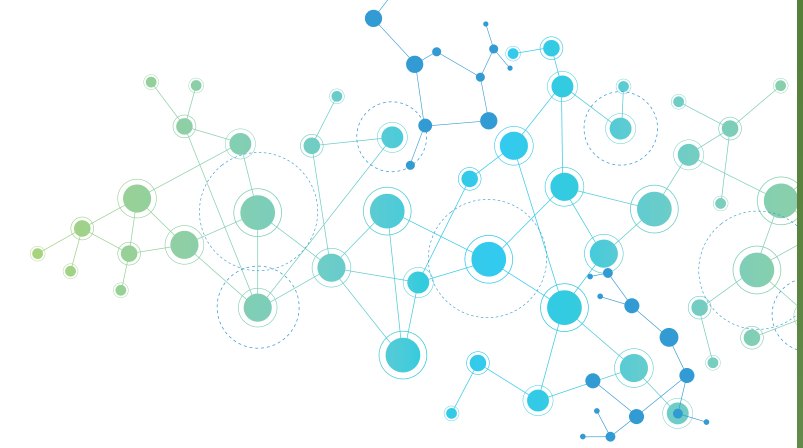Research Professor of Computer Science

An aerial view of Marina del Rey, California, with the Pacific Ocean and the Santa Monica Mountains in the background, and Information Sciences Institute in the foreground (13–story building at lower left).

ISI's East Coast offices are located in Arlington, VA, in close proximity to Washington D.C. and Northern Virginia's technology centers.

The title of Research Director was established by ISI to recognize our leaders who head a sizable scientific research group, and develop, plan and direct the research and funding strategy for that group. All of our Research Directors make critically valuable contributions to the current and future success of the Institute. We are pleased to introduce the following researchers who are joining existing directors Yolanda Gill and Craig Knoblock as ISI Research Directors.

Stephen Crago

Director of the Computational Systems & Technology Division and a research associate professor in EE, Stephen is a leader in heterogeneous cloud and embedded computing. He served as deputy director of CS&T from 2013-2015, and is currently the director of ISI's Arlington campus. Stephen publishes extensively and is active in the professional community. His leadership contributions include service as a member of ISI's Infrastructure Oversight Committee, and leadership of the IC business development group and the Arlington security team.

Matthew French

As a senior computer scientist and a recognized leader in reconfigurable computing, trusted electronics, and wireless technology, Matthew contributes significantly to the leadership of ISI's Computational Systems & Technology Division and to its Arlington campus. He is responsible for significant collaboration — both internally (across ISI's Arlington and Marina del Rey campuses) and externally (with top universities in reconfigurable computing). He is the founding director of the SURE Center at ISI, composed of a team of technical leaders from the trusted electronics areas within ISI.
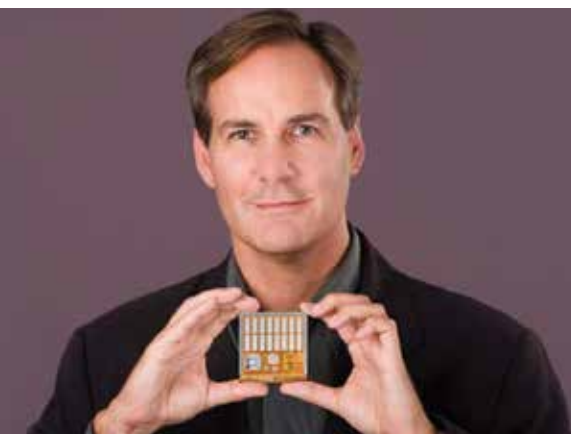
Ewa Deelman

As a research professor and pioneer in the area of science automation, Ewa's work has received over 13,000 citations. Her Pegasus workflow tool is widely used in astronomy, bioinformatics, earthquake science, gravitational wave physics, ocean science, limnology, and in other scientific communities. Pegasus was recently used in one of the main analysis pipelines in the LIGO experiment. In 2015, Ewa received the HPDC Achievement Award for making long-lasting, influential contributions to the foundations or practice of the field of high-performance parallel and distributed computing (HPDC).

Aram Galstyan

A research associate professor in CS who joined ISI in 2000, Aram has not only distinguished himself as a researcher in data science and machine learning, he has also been instrumental in recruiting superb research talent to ISI. Whether by himself, in collaboration with others within his group, or through working with other veteran ISIers, Aram has made significant contributions in the areas of machine learning, computational social science, and bioinformatics.

Jeffrey Draper

A Deputy Director of the Computational Systems & Technology Division and a research associate professor in EE, Jeffrey has a proven track record as a researcher and leader in the areas of VLSI design, computer architecture, and trusted electronics. He is an active author, Ph.D. advisor and educator, and contributes to faculty leadership through the Viterbi School of Engineering's Faculty Council. Jeff is also an Assistant Director of the newly established SURE (Secure and Robust Electronics) Center.

Michael Orosz

In his roles as research associate professor and technical leader, Michael's focus is that of applying information sciences to solve complex problems in mission-critical operational environments. He joined ISI in 2001, and the following year became leader of the Decision Systems group. In addition to his ISI appointments, he is a research associate professor in the Sonny Astani Department of Civil & Environmental Engineering and is currently interim executive director for Corporate and Foundation Relations in Viterbi's Advancement group. In spite of these demands, Michael continues to be active in academic and sponsored committees and panels.

NEW ISI RESEARCH LEADS

The title of Research Lead was established by ISI to acknowledge researchers who continuously make noteworthy contributions to their field of research and to Information Sciences Institute. We are pleased to announce the following researchers who have joined the ranks of our existing ISI Research Leads: Gully Burns, Hans Chalupsky, Yu-Han Chang, Andrew Schmidt, Federico Spedalieri, Mehdi Yahyanejad, and Ke-Thia Yao.

Wael AbdAlmageed

Wael joined ISI in 2013 as a research scientist in the then nascent computer vision research group. Prior to joining ISI, he was a member of the research staff at the University of Maryland's Institute for Advanced Computer Studies. Since joining ISI, Wael has made significant contributions to several proposals, including two that resulted in major research awards from IARPA and DARPA under their Janus and MediFor programs respectively. His contributions have been important in establishing a vibrant and impactful computer vision research group at ISI.

Greg Ver Steeg

Greg joined ISI in 2009 after receiving his doctorate in physics from Caltech. As a research lead and research assistant professor, his work focuses on using ideas from information theory to understand complex systems like human behavior, biology, and language. One area of this research targets abstract representations that help us understand complex data like the neurophysiology of Alzheimer's patients and gene expression of cancer patients. In recognition of his achievements, Greg received a *Young Investigator Award* from the Air Force Office of Scientific Research for 2012-2015.

INFORMATION SCIENCES INSTITUTE ACHIEVEMENT AWARD

In October 2015, the scientific and administrative staff at Information Sciences Institute gathered together to acknowledge excellence and leadership in research by awarding Terry Benzel and Kevin Knight the inaugural Institute Achievement Award. Terry Benzel is ISI's deputy division director for the Internet and Networked Systems Division, and Kevin Knight is the director of the Natural Language Technologies Group.

Terry Benzel is also the technical project lead for the Cyber Defense Technology Experimental Research (DETER) testbed projects funded by DHS, NSF and DARPA. The projects are developing an experimental infrastructure network and rigorously testing frameworks and methodologies to support the development and demonstration of next-generation information security technologies for cyber defense. She is also an advisor to government and industry on R&D strategy and holds a joint appointment at the Institute for Critical Information Infrastructure Protection, part of the USC Marshall School of Business.

Terry Benzel

Institute Achievement Award for "advancing ISI's reputation by creating and leading a vibrant, sustained, and influential program of research in experimental cybersecurity methodology and infrastructure."

Kevin Knight is a renowned leader is the area of intelligent systems and natural language processing. He is currently leading the new *ELISA* project (under DARPA's LORELEI program) which anticipates a natural or man-made disaster somewhere in the world that would naturally precipitate a flurry of Internet communication and news in the local language—which would be just one of the world's 4000+ languages. ELISA will develop technologies to quickly field NLP tools for that local language, translating text, extracting important entities and their relationships, and detecting sentiment/emotion in both text and speech with the overall goal of providing direction for humanitarian assistance.

Kevin Knight

Institute Achievement Award for "leadership in recruiting future generations of researchers to ISI."

Scientists and Engineers

Tameem Albash	Senior Research Associate
Huaigu Cao	Computer Scientist
Joshua Chudy	Programmer Analyst II
Emilio Ferrara	Computer Scientist
Tal Hassner	Senior Computer Scientist
Anoop Kumar	Senior Computer Scientist
Hongsuda Tangmunarunkit	Computer Scientist
Gabriel Weisz	Computer Scientist
Jessie Wong	Programmer Analyst I
Barret Zoph	Research Programmer I

Administrative Staff

Lily Fierro	Business Data Analyst
Derek Lazzaro	Chief Information Officer
Karen Rawlins	Executive Secretary
Janice Wheeler	Technical Grant Writer

Postdoctoral Scholars

Yonatan Bisk	Natural language processing, computational linguistics
Daniel Garijo	e-Science and the semantic web, artificial intelligence
Mohsen Taheriyan	Information integration, semantic models of data sources
Xiaoran Yan	Social network dynamics and community detection

Ph.D. Graduates

Ph.D. Graduates	Advisor	Research	Currently at:
Hashem Alayed	C. Neuman	Security in online games	King Saud Univ. (Faculty)
Siddharth Bhargav	Y. Cho	Highly accurate on-line on-chip subcircuit power monitoring	Nvidia Inc. (ASIC Power Engr)
Qing Dou	K. Knight	Natural language processing	Facebook (Research Scientist)
Xun Fan	J. Heidemann*	Networking	Microsoft (Technical Staff)
Daniel Garijo	Y. Gil	Mining abstractions in scientific workflows	USC/ISI (Postdoc)
Andrew Goodney	Y. Cho	Micro ocean acoustic tomography/wireless sensor NW	USC CS Dept. (Lecturer)
David Lee	J. Draper	Reliability of highly scaled field-programmable gate arrays	Sandia Nat'l Labs (Res. Staff)
George Konstantinidis	J. L. Ambite	Data integration/exchange	Oxford Univ. (Postdoc)
Gopi Neela	J. Draper	Optimizations for 3D integrated circuits	Nvidia Inc. (Engineer)
Mohsen Taheriyan	C. Knoblock	Learning semantics of structured data sources	Google (Software Engineer)
Bo Wu	C. Knoblock	Data integration & transformation, AI and user interfaces	Yelp (Software Engineer)
Arun Viswanathan	C. Neuman	Situational awareness & intrusion detection	JPL (Research Technologist)

Co-advisors Ramesh Govindan and Ethan-Katz Basset

Doctoral Students

Doctoral Students	Research Area	Advisor
Abdulla Alwabel	Computer systems security, distributed systems	Clifford Neuman
Calvin Ardi	Framework for sharing network security data	John Heidemann
Luenin Barrios	Autonomous robots and systems	Wei-Min Shen
Joseph Chen	Autonomous robots and systems	Wei-Min Shen
Yu-An Chen	Heterogeneous real-time cloud computing	Stephen Crago
Thomas Collins	Autonomous robots and systems	Wei-Min Shen
Xiyue Deng	Network security defense	Jelena Mirkovic
Aliya Deri	Natural language technologies	Kevin Knight
Aditya Deshpande	Memory architecture, computer arithmetic	Jeffrey Draper
Abdessamad Echihabi	NLP focus on conversational understanding	Daniel Marcu
Shuyang Gao	Information theory, machine learning	Aram Galstyan
Sahil Garg	Data-driven modeling of complex systems	Aram Galstyan
Majid Ghasemi Goi	Data link discovery—DIG project	Pedro Szekely
Marjan Ghazvininejad	Natural language processing	Kevin Knight
Hang Guo	Analysis of network traffic	John Heidemann
Shawn Kailath	Self sustaining cyber-physical systems	Young Cho
David Kale	Machine learning in high-impact domains, incl. health care	Greg Ver Steeg
Andrew Kolb	Biomimetic audio signal processing	Young Cho
Farshad Kooti	Large, complex datasets, especially online social networks	Kristina Lerman
Neal Lawton	Applied machine learning	Aram Galstyan & Greg Ver Steeg
Ji Li	Computer architecture	Jeffrey Draper
Kuan Liu	Machine learning and natural language processing	Prem Natarajan
Daniel Moyer	Machine learning, connectomics, network theory	Greg Ver Steeg
Minh Pham	Semantic labeling, data integration, machine learning	Craig Knoblock
Nima Pourdamghani	Natural language technologies	Kevin Knight
Abdul Qadeer	Distributed sys & operating sys; networks & architectures	John Heidemann
Kyle Reing	Statistical machine learning & complex networks	Aram Galstyan & Greg Ver Steeg
Ekraam Sabir	Optical character recognition	Prem Natarajan
Praveen Sharma	Computer architecture	Jeffrey Draper
Vinod Sharma	Network security defense	Jelena Mirkovic
Emily Sheng	Natural language processing	Prem Natarajan
Hao Shi	Network security defense	Jelena Mirkovic
Xing Shi	Natural language technologies	Kevin Knight
Jason Slepicka	Databases, artificial intelligence, information sciences	Pedro Szekely
Geoffrey Tran	Computer architecture, multiprocessors, cloud computing	Stephen Crago
Yatin Wadhawan	Healthcare data security/cyber-physical systems	Clifford Neuman
Lan Wei	Internet measurement of IP anycast	John Heidemann
Simon Woo	Network security defense	Jelena Mirkovic
Hao Wu	Analysis of network traffic	Kristina Lerman
Xin-Zeng Wu	Social dynamics and networks	Kristina Lerman
Liang Zhu	Domain name system and analysis	John Heidemann

Master's & Undergraduate Researchers

Suresh Alse Aniket

RajaGopal Bojanapalli

Yinyi Chen

Ayesha Dastagiri

Snehal Surendra Desai

Nikunj Gala

M. R. Gurumoorthy

Andy Gu

Rebecca Hao

Yueheng He

Jiayuan Huang

Mandar Kapdoskar

Jyothsna Kullatira

Geoffrey Knopf

Neha Reddy Lakkadi

Vishnukarthik Lakshmanan

Stephen Lee

Mengjia Liang

Sanmukh Sanjog Lodha

Weiqiang Luo

Shuai Ma

Fangxu Meng

Saurabh Mishra

Matthew O'Brien

Vidyadhar Pednekar

Sharanya Radhamohan

S. Satyamangalam Ramanathan

Jaydeep Ramani

Satheesh Prabu Ravi

Gaurav Sudhir Sathe

Aref Shafaeibejestan

Kinara Himanshu Shah

Jayavanth Shenoy

Preetam Shingavi

Anurag Shukla

Dimitrios Stripelis

Zheng Tang

Ankita Tiwari

Kishan Trivedi

Swaroop Venu

Sowmya Naidu Vundamati

Wenqiang Wang

Renjie Zhao

Yuhao Zhu

Visting Scholars & Researchers

Emanuel Alsina, Research Scholar & Ph.D. candidate
University Modena & Reggio Emilia, Italy
Diversity/complexity in group dynamics and its impact on performance

Ricardo De Oliveira Schmidt, Research Scholar
University of Twente, Netherlands
Network traffic measurements and analysis

Luciano Gallegos, ISI Postdoctoral Fellow
Universite de Rennes, France
Analysis of emotions in social media; census data using machine learning and statistical data

Gleb Gawriljuk, Short-Term Scholar & M.S. student
Karlsruhe Institute of Technology, Germany
Intelligent systems research

Gerald Hiebel, Research Scholar & Ph.D.
University of Innsbruck, Austria
Integration and modeling of geospatial data using semantic technologies and ontologies

Patricia Jumenez Aguirre, Short-Term Scholar & M.S. student
Universidad de Sevilla, Spain
Internet information extraction technology

Sung Dong Kim, Research Scholar & Ph.D.
Seoul National University, Korea
Development of English-Korean machine translation systems

Juliane Kruck, Research Scholar
Karlsruhe Institute of Technology, Germany
Guarantees of valid scientific results and related privacy issues

Viviana Natase, Short-Term Scholar & Ph.D. Researcher
Fondazione Bruno Kessler, Italy
Intelligent systems research

David Oaknin, Research Scholar & Ph.D.
RAFAEL Advanced Defense Systems, Israel
Quantum physics

Fernando Ortega Requena, Short-Term Scholar & Ph.D. candidate
Technical University, Madrid, Spain
Techniques for handling Big Data

Jan Pesan, Research Scholar & Ph.D. candidate
Brno University of Technology, Czech Republic
Artificial intelligence, algorithms, computer security

Longhua Qian, Visiting Researcher & Ph.D.
Soochow University, China
Information extraction, text mining, bilingual resource processing, social networks

Qiang Li, Research Scholar & Ph.D. Candidate
Northeastern University, China
Natural language processing and machine translation

Jing Qiu, Visiting Researcher
Hebel University of Science and Technology, China
Information extraction in ontology learning

Sudha Rao, Research Scholar
University of Maryland
Natural language processing

Stephan Schlagkamp, Research Scholar & Ph.D. candidate
TU Dortmund University, Germany
Optimization goals and performance evaluation in high-performance computing

Philipp Singer, Short-Term Scholar & Postdoc Researcher
Institute for Social Sciences, Germany
Data science, statistics, machine learning, web science

Dishelt Torres Paz, Short-Term Scholar & M.S. student
Trier University, Germany
Replicated data management architecture

Wenduan Xu, Short-Term Scholar & Ph.D. candidate
University of Cambridge, UK
Artificial neural networks, algorithms

Hanbing Yao, Research Scholar & Ph.D.
Huazhong University of Science and Technology
Web data mining, text sentiment analysis, information retrieval

Deniz Yuret, Research Scholar & Ph.D.
Koc University, Istanbul, Turkey
Machine vision, natural language processing, computational linguistics, machine vision

Awards

Genevieve Bartlett and **Jelena Mirkovic**
Best Paper Award at Int'l Workshop on Computer and Networking Experimental Research using Testbeds (CNERT) for "Expressing Different Traffic Models Using the LegoTG Framework," 2015.

Terry Benzel
ISI Institute Achievement Award, 2015.

Ewa Deelman
2015 ACM HPDC Award for long-lasting, influential contributions to the foundations or practice of the field of high-performance parallel and distributed computing.

Emilio Ferrara
Named *VIP Influencer* in the field of big data by IBM Watson Analytics and PureMatter | Elected *Senior Fellow* of UCLA's Institute for Pure and Applied Mathematics | *Top 100 Big Data Experts* selected by Maptive | *Best Poster Award* for "The Rise of Social Bots: Bot Detection in Social Media," 2015 Conference on Complex Systems.

Itay Hen
DOE's Advanced Scientific Computing Research (ASCR) *Leadership Computing Challenge* awards for 2014-2015 and 2015-2016.

Kevin Knight
Fellow, Association for Computational Linguistics (ACL), 2015
ISI Institute Achievement Award, 2015.

Daniel Marcu
Fellow, Association of Computational Linguistics (ACL), 2015.

Pedro Szekely, Craig Knoblock, Jason Slepicka, Andrew Philpot, Amandeep Dingh, Chengye Yin, Dipsy Kapoor, Prem Natarajan, Daniel Marcu, Kevin Knight, David Stallard, Subessware Karunamoorthy, RajaGopal Bojanapalli, et al.
Best Applied Paper Award at 14th International Semantic Web Conference (ISWC) for "Building and Using a Knowledge Graph to Combat Human Trafficking," 2015.

Gabriel Weisz and James Hoe
Best Paper Award at 25th International Conference on Field Programmable Logic and Applications (FPL 2015) for "CoRAM++: Supporting Data-Structure-Specific Memory Interfaces for FPGA Computing," 2015.

Activities

Wael AbdAlmageed
Program Committee, Machine Learning in High Performance Computing Environments, SuperComputing 2015.

José Luis Ambite
Program Committee Chair, Int'l Conference on Data Integration in the Life Sciences (DILS2015).

Yigal Arens
Program Committee: IEEE/ACM ASONAM Workshop on Intelligence and Security Informatics 2015, 2015 Int'l Conferences on Digital Government Research, Int'l Conference for Smart Health (ICSH) 2015 | *Co-Chair*, Smart Cities Track, 2015 Int'l Conferences on Digital Government Research.

Terry Benzel
Associate Editor-In-Chief, IEEE Security and Privacy Magazine, 2015-2016 | *Organizer*, GREPSEC II Workshop for Women and Underrepresented Minorities in Cyber Security, 2015.

Gully Burns
General Reviewer for Frontiers in Neuroinformatics and *Guest Editor* for special topic: "Discovery Informatics in Neuroscience Extracting and Shaping Actionable Knowledge from Texts and Data."

Young Cho
Technical Program Committee ACM/IEEE ANCS 2015.

Stephen Crago
Steering Committee Chair, IEEE Transactions on Big Data
Area Chair, IEEE Cloud 2016, Systems Software and Hardware
Organizing Committee, Fault-Tolerant Spaceborne Computing Employing New Technologies.

Ewa Deelman
2015 Editorial Board Member, Int'l Journal of High Performance Computing Application and ACM Int'l Conference Proceeding Series | *Co-Technical Papers Chair*, SC'15 Technical Program
Vice-Chair of Program Committee, 11th Int'l Conference on Parallel Processing & Applied Math | *Tutorials Co-Chair*, 7th IEEE/ACM Int'l Conference on Utility & Cloud Computing, 2015.

Pedro Diniz
General Co-Chair of 2015 Applied Reconfigurable Computing Symposium (ARC15) | *Technical Program Committee:* DATE'15, SAMOS'15, and FPL'15.

Activities (cont.)

Matthew French
IEEE Senior Member | *General Chair* and *Technical Program Committee Member*, 2016 Int'l Symposium on Field Programmable Custom Computing Machines.

Emilio Ferrara
Program Chair of 8th Int'l Social Informatics Conference
Co-Chair of CCS 2016 workshop | *Invited Talks:* Northwestern University (predicting human behaviors in techno-social systems), Stanford Univ. (predictability of rare events leveraging social media, a machine learning perspective), and two talks at the Conference on Complex Systems 2015 | *USC Guest Lecturer* for Special Topics/Social Networks and for Introduction to Computational Thinking and Data Science | *Editorial Board* for Data Science (IOS Publishing) and *Guest Editor* of Computational Social Sciences special issue on Future Internet, 2015-2016 | *Conference Program Committee:* IJCAI 2016, ACM Web Science 2016, ICWSM 2016, ACM WWW 2016, ASONAM 2015, SocInfo 2015, SocialCom | *Workshop Program Committee:* CoupleNet 2016, CSSWS 2015, SidEWayS 2015, #Fail! 2015.

Yolanda Gil
General Chair, Int'l Semantic Web Conference (ISWC) 2016
Chair/Co-Chair: ACM's Special Interest Group In Artificial Intelligence (SIGAI), Technology and Architecture Committee of the NSF EarthCube program, and Workshop on Scientific Knowledge Capture at ACM Int'l Conference on Knowledge Capture, 2015 | *Founding Board Member*, Leadership Council of NSF EarthCube program | *Advisory Board Member*, Computer Science Section, MethodsX Journal, and Semantic Web Science Association *Founding Editorial Board Member*: Journal Data Science–Methods, Infrastructure and Applications, and AI Matters (ACM SIGAI Quarterly Newsletter) | *Editorial Board:* Artificial Intelligence Journal, Journal of Web Semantics, and Applied Ontology Journal.

Itay Hen
Invited Talk, Fidelity-optimized quantum state estimation, Quantum Machine Learning Workshop (NIPS), 2015.

Kevin Knight
General Chair, Conference of the North American Association for Computational Linguistics (NAACL HLT), 2016 | *Program Committee,* ACL 2015, EMNLP 2015 | *Student Research Workshop faculty advisor*, ACL 2015 | *Co-Organizer*, Dagstuhl Workshop on "Graph Transformation for Natural Language Processing" | *Invited Talks*: Nodalida Conference in Vilnius, Lithuania 2015, ACL Bio-NLP Workshop 2015, and Midwest Language Workshop 2015.

Craig Knoblock
President, Int'l Joint Conferences on AI, 2013-2015 | *Trustee*, Int'l Joint Conference on AI, 2007-2015 | *Senior Program Committee*, Conference on AI (AI & the Web track), 2015
Program Committees: Workshop on Linked Data for Information Extraction 2015, 11th Int'l Conference on Data Integration in the Life Sciences 2015, and ISWC 5th Int'l Workshop Linked Science 2015 | *Associate Editor*, Artificial intelligence 2009-2016 and Spatial Algorithms and Systems 2015-2016 | *Editorial Board,* Journal of Web Semantics 2010-2015 and AAAI Press 2015
Advisory Board, ACM Transactions on Intelligent Systems & Technology 2015-2016.

Daniel Marcu
Advisory Board of Journal of Artificial Intelligence Research
Associate Editor of Artificial Intelligence Journal. *Associate Editor* | *JAIR Advisory Board Member* | *Program Committee:* ACL-IJCNLP 2015, IJCAI 2015, NAACL 2015, EMNLP 2015, AAAI 2016, MT Summit 2015, WPTP 2015, ACL Student Workshop 2015, ICCP 2015

Clifford Neuman
Program Committee, 2015 Industrial Control Systems Security Workshop at ACSAC and 2015 IEEE Int'l Conference on Intelligence and Security Informatics | *Editorial Board* for Int'l Journal of Electronic Commerce.

Andrew Schmidt
Co-Chair for High Performance Computing Track, Int'l Conference on Reconfigurable Computing and FPGAs
Technical Program Committee Member, Int'l Symposium on Field Programmable Custom Computing Machines.

Stephen Schwab
General Chair, Annual Computer Security Applications Conference, 2015.

Joseph Touch
Presentations: OSA Frontiers in Optics-Nonlinear Optics Tech Group Workshop, 2015; OSA Optical Computing Incubator Workshop, 2015; George Washington Univ. Electrical Engineering, 2016; MIT Lincoln Laboratories, 2016; Boston Univ. Computer Science, 2016, Stanford Univ. Electrical Engineering, 2016; UCSD Electrical Engineering, 2016; HP Laboratories, 2016.

The Secure and Robust Electronics Center (SURE) | *Director Matthew French*

Global production trends have introduced significant vulnerabilities into electronics hardware. Integrated circuits that once were created by a single company, possibly under a single roof, are now produced by "fabless foundries" with few or no facilities of their own. Instead, state-of-the-art production relies on a complex, international design and manufacturing process. The same development flow and novel physics that are pushing chips to new levels also expose new security risks.

Designs may be generated with tools from multiple software vendors. Source code may be written by different companies, possibly in different countries. Key technical elements may be licensed from third parties. Foundries and their customers rarely know with absolute certainty whose hands a chip has passed through, whether a latent issue has occurred, and how the chip will perform in all circumstances.

In addition, modern nanoscale fabrication is adding complexity to reliability and resiliency issues. The performance of individual transistors has become increasingly variable, making overall system reliability challenging. Shrinking voltage margins increase susceptibility to errors that historically were seen only in space environments. Additionally, wear-out and aging effects appear sooner in the lifetime of integrated circuits and with more variability.
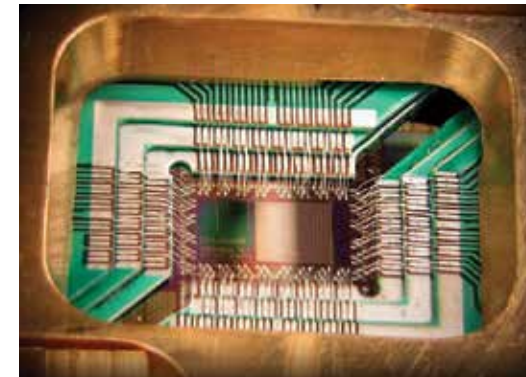
The issues could have profound consequences for everything from misfires or failures of major U.S. infrastructures, both physical and virtual, to individual smartphone performance.

USC's Information Sciences Institute established the *Secure and Robust Electronics Center (SURE)* to address these issues. SURE seeks to investigate and address state-of-the-art manufacturing complexities that have significantly compromised four key aspects of chip production: trust, security, resilience and reliability. SURE researchers have over a decade of experience in this emerging field, with supporters that include DARPA, IARPA, NASA, Army Research Office, Air Force Research Laboratory and Defense Threat Reduction Agency. As experienced industry collaborators, we aim to accelerate the scale, pace and impact of hardware robustness and security technology development; that which were once military and aerospace concerns, are now rapidly expanding into the commercial sector.

The SURE Center combines the research thrusts of ISI with MOSIS's fabless foundry capabilities, enabling SURE to address full system device design and production, including front-end design and FPGA programming, microarchitecture, integrated circuits, CAD, and fabrication.

USC-Lockheed Martin Quantum Computation Center | *Directors Federico Spedalieri, Daniel Lidar*

The USC-Lockheed Martin Quantum Computation Center (QCC) is a ground-breaking joint scientific research effort between the Lockheed Martin Corporation and the USC ISI. Jointly operated by ISI and Lockheed Martin, the QCC seeks to bridge the gap between academic research in quantum computation and critical real-world problems.

Faculty, researchers and students are performing basic and applied research into quantum computing, and are collaborating with researchers around the world. The QCC houses a D-Wave Two quantum annealing system, manufactured by D-Wave Systems, Inc. The QCC is the first organization outside of D-Wave to house and operate its own system. Several years after the installation at USC ISI, a second system was installed at NASA Ames Research Center and is operated jointly by NASA and Google. Operating quantum computing systems is demanding: the systems need to be kept near absolute zero temperature and electromagnetically shielded to maximize possible coherence time.
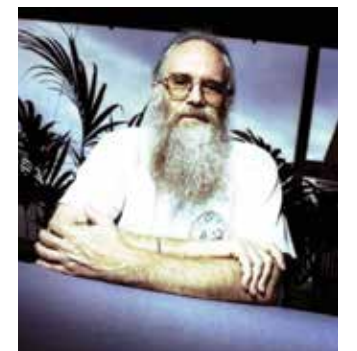
Quantum computers use quantum bits, or qubits, which can encode a one and a zero at the same time (Photo courtesy of D-Wave Systems Inc.)

The Postel Center for Experimental Networking | *Director Joe Touch*

The Postel Center for Experimental Networking is an endowed facility focusing on network research and service to the Internet community. In its early years, the Center focused on Internet history and supporting network research tools and resources through visiting scholars. Recently, we have explored Internet policy issues and issues underlying our understanding of the Internet architecture and its protocols. We developed reference documents for the Internet community to clarify architectural concepts (tunnels, middleboxes, GRE fragmentation, transport port use) and to extend legacy protocols for new uses (both TCP and UDP option space extensions). Today's mission for the Postel Center is to promote Jon's tradition of applied research service to the internet community and to explore new missions and future directions.

The Postel Center is named in memory of **Dr. Jon Postel**—a brilliant and dedicated scientist who made many key contributions to the formative days of the ARPANET, including protocol design and verification, multimedia computing and communications, electronic commerce, the domain name system, and many other specific Internet protocols. Jon was widely known for the influence he exerted on the management of the Internet, recognizing early on that packet-switching research would need organization and a modicum of discipline if it were to realize its full potential as a universal communication medium. Jon created the activities that eventually grew into the RFC Editor, which issues and controls the many documents that specify how Internet computers interoperate, and also initiated (circa 1981) the Internet Assigned Numbers Authority (IANA), the central coordination function for the global Internet.

Dr. Jon Postel, key contributor to the formative days of the Internet

The Center for Computer Systems Security | *Director Clifford Neuman*

The ISI *Center for Computer Systems Security* conducts research and provides education in the crucial disciplines of computer, network, and application security. Among the current research activities, Center staff have studied resilience to cyber-attack in the critical infrastructure systems of the power grid and oil and gas extraction.

This work demonstrates how cyber-attacks can affect the operational resilience of the infrastructure, impeding the delivery of power to consumers or oil and gas to refineries. It also identifies the most effective remediation strategies among candidate actions, such as system reconfiguration and acquisition of a reserve capacity that will maximize the resilience metric for the system that is under attack.

Center researchers are currently exploring cross-infrastructural dependencies to identify how loss of storage or transmission capacity in one infrastructure impacts the resilience of other infrastructures in the same manner that loss of gas storage in the Alisio Canyon storage field in Los Angeles will affect the resilience of the power grid in summer 2016.

In addition to its research activities, Center researchers are frequently called upon by the media to explain events involving privacy, cyber security, and cyber crime. In the past year, Center Director Clifford Neuman was quoted in more than 40 publications and appeared more than 15 times in television and national network news segments and radio programs, where topics ranged from the Apple vs. FBI case, to hospital ransomware, and privacy and security for network-connected devices and vehicles.

ISI's Center for Computer Systems Security is a DHS and NSA designated center of academic excellence in information assurance research.

Center staff lead several education programs in computer security and are redesigning the master's of science degree in security engineering. They are adding a cross-disciplinary class in privacy and are restructuring many classes in order to apply the existing focus of the program in high assurance systems to the Internet and cloud-focused computing environments now in demand by consumers and security.

Leidos—Research on Human-Computer Interaction

Leidos, Inc., a national security, health, and engineering research firm headquartered in Arlington, VA, and Information Sciences Institute are collaborating to provide scientific, engineering, systems integration and technical services to the U. S. Defense and Intelligence communities, civil agencies, and selected commercial markets. Leidos and researchers at Information Sciences Institute are currently working together on the following significant research efforts.

ComplexGraphs

Networks capture complex relationships between entities (nodes), whether they are people interacting with social networks or entities, events and places within a knowledge graph. Existing network analysis can quantify network structure to answer questions, such as—Which are the important, or central nodes within the network? What hidden groups or communities exist? Which new interactions are most likely to be observed in the future? Computationally efficient algorithmic solutions to these and related questions have made a tremendous impact—as witnessed, for example, by the success of Google's PageRank algorithm. Other equally important questions, however, have been more difficult for existing network analysis tools to answer. These questions include—How similar are the two networks? What nodes play similar roles in the same network? What nodes play similar roles in different networks? Inspired by recent success of "deep learning" and neural language models for text analysis, *ComplexGraphs* is developing new algorithms to answer these questions.

Evaluation of Mapping NLP-Based Algorithms into Reconfigurable Computing Kernels (EMBARK)

Human-computer interaction has evolved significantly with the introduction of digital assistants such as Siri, Cortana, and Alexa, opening the door to what next-generation computing platforms will be capable of achieving. At the heart of these technologies is natural language processing, which allows a computer to extract meaning from natural language input instead of structured computer code, such as C++ or Java. The complex algorithms developed to enable such capabilities as information extraction, information retrieval, and sentiment analysis are pushing the envelop of existing computing platforms. CPU- and GPU-based systems have typically been leveraged for their ease of programmability; however, these systems still have performance and power limitations. *EMBARK* looks towards reconfigurable computing using field programmable gate arrays (FPGAs) which allow for low-level hardware customizations for each application—this is akin to swapping out specialized processors for each task. With FPGA-based systems it is possible to not only construct application-specific computation accelerators, but to tightly integrate them within efficient memory and network interfaces to better support scalability and meet the demands of the evolving algorithms. EMBARK looks to analyze these algorithms to develop a framework to efficiently support decomposing the building blocks of the algorithms into configurable computing kernels and interfacing these kernels in FPGA-based heterogeneous systems, overcoming inherent FPGA limitations for slow clock speeds, programmability, and run-time management.

The Laser Interferometer Gravitational-Wave Observatory (LIGO)

Aerial view of the LIGO observatory at Hanford, Washington state (Credit: Caltech/MIT/LIGO Laboratory)

Jointly operated by Caltech and MIT, supported by the National Science Foundation and Department of Energy, and an international resource for both physics and astrophysics, the *Laser Interferometer Gravitational-Wave Observatory* was built upon the early work of many scientists to test a component of Albert Einstein's theory of relativity — the existence of gravitational waves (predicted by Einstein in 1916). Beginning in the 1960s, American and Soviet scientists conceived the basic ideas and prototypes of laser interferometry, and In 1968 Caltech initiated theoretical efforts on gravitational waves and their sources. Since that time, LIGO has been a true collaborative effort, involving more than 1000 scientists from dozens of institutions and 16 countries worldwide. The *Pegasus Project* at USC's Information Sciences Institute (https://pegasus.isi.edu) has been involved with LIGO since 2001, when both groups became part of a National Science Foundation-funded project to explore the use of workflow technologies to run large-scale science pipelines on computational grids. Pegasus enabled the analysis of data gathered from the twin LIGO detectors, located in Louisiana and Washington state, and the GEO600 detector in Hannover, Germany. Initially, pipeline analyses were managed by the Pegasus Workflow Management System on the LIGO Data Grid. LIGO then extended its computations to nationwide cyber-infrastructures, such as NSF's Open Science Grid (OSG) and Extreme Science and Engineering Discovery Environment (XSEDE). Pegasus contributed to this expansion by managing cross-site data transfers over wide-area networks, and large-scale computations in a reliable, scalable, and efficient manner.

The benefits of using Pegasus to run LIGO analyses are many: 1) *an expanded computing horizon*—Pegasus enabled LIGO pipelines to run across external cyberinfrastructures beyond the LIGO collaboration, which consequently leveraged other research infrastructures funded by the National Science Foundation (e.g., OSG and XSEDE) and European collaborations (VIRGO); and 2) *smart data management*—Pegasus' optimization strategies enabled automated data discovery (e.g., symbolic links to locally available input data, fallback to remote file servers if data is not available locally, support for retrieving data using various protocols, etc.), automated data cleanup to reduce peak storage requirements, data reuse to reduce the amount of storage required for the computation, and support for checkpointing.

Early this year, physicists at the Laser Interferometer Gravitational-Wave Observatory announced a long-awaited scientific breakthrough: the detection of gravitational waves as predicted by Einstein's Theory of Relativity. USC's Information Sciences Institute is proud of its association with LIGO and its contributions to LIGO's outstanding scientific achievements.

Northrop Grumman Cybersecurity Research Consortium

Northrop Grumman's Cybersecurity Research Consortium, founded in 2009, is a groundbreaking partnership of industry and academia formed to advance research, facilitate collaboration among the nation's top scientists, and accelerate solutions to counter the fast-changing threats from cyberspace. The consortium addresses some of the world's leading cyber problems, including attribution in cyberspace, supply chain risk, and securing critical infrastructure networks. Members of the consortium coordinate research projects, share information and best practices, develop curricula, write joint case studies, and provide numerous learning opportunities and applications for students and the defense community overall.

USC's Viterbi School of Engineering joined the Northrop Grumman Cybersecurity Research Consortium in 2013, with I*nformation Sciences Institute* serving as the lead organization for establishing and advancing the partnership and expanding the consortiums's breadth of investigation into the most pressing cyber threats to the economy and national security. ISI brought with it a strong reputation for leadership in big data, cybersecurity, computer science, and informatics. The other three partners in this research consortium are also leading cybersecurity research universities—Carnegie Mellon University, Massachusetts Institute of Technology, and Purdue University.

ISI researchers are collaborating with researchers at Viterbi's Energy Institute and its Center for Energy Informatics on developing models to understand the impact of attacks and failures in the power grid using situational awareness of both the power grid and gas distribution infrastructures. This will allow a greater understanding of the resilience of the power grid under conditions such as the recent shutdown of the Aliso Canyon's underground gas storage facility which supplies natural gas to "peaker" plants (that generally only run when there is a "peak" demand) operated by LA's Department of Water and Power and Southern California Edison.

Collaborative research aimed at securing critical infrastructure networks.

Computational Systems and Technology Division | *Director Stephen Crago*

ISI is one of the world's leaders in computing technologies, ranging from basic and applied research to hands-on development. Our Computational Systems and Technology (CS&T) Division focuses on:

Compilers and big data

Decision systems

Heterogeneous cloud and embedded computing

Microarchitecture, integrated circuits and advanced electronics

Quantum computing

Reconfigurable computing and wireless networks

Science automation technologies

Trusted electronics and computing

From Theoretical to Hands-On

Our current initiatives include theoretical adiabatic quantum computing through the **USC-Lockheed Martin Quantum Computation Center** and hardware security through the **Secure and Robust Electronics Center**. CS&T projects include high-performance compilers and software for heterogeneous clouds, and hardware-software design of unique chips and field-programmable gate arrays. We're also exploring algorithms and data structures to help understand and control large-scale, complex systems such as oil field geometry, megacity emergency response and legacy software security.

CS&T teams are creating wireless networking technologies for battlefields and other difficult environments, along with social media for people who lack trustworthy Internet access. We're advancing our scientific automation tools—which enable researchers to focus on conducting science, not managing data—already used by astronomers, physicists and earthquake specialists, including the LIGO team that was the first to detect gravitational waves that Einstein predicted. Our team also is deeply involved in the Los Angeles Department of Water and Power's citywide smart grid project.

Intellectual Leadership, Diverse Expertise

We provide intellectual leadership within the broader research community through participation in workshops like the White House OSTP **National Strategic Computing Initiative Workshop** and efforts such as the **Institute for Sustained Performance, Energy and Resilience (SUPER)** supercomputing project, a group of four national laboratories and eight universities seeking to alleviate the tedious burden of importing large-scale code from one computing system to another. The division's 40 computing technology researchers, research programmers and graduate students represent a wide range of disciplines, including electrical engineering, computer science, physics, math, and psychology. CS&T researchers are based in Arlington, Virginia and in Marina del Rey, California.

Informatics Systems Research Division | *Director Carl Kesselman*

ISI's Informatics Systems Research Division pursues a broad research agenda focused on creating new types of sociotechnical systems that enable and accelerate discovery in domains of high societal impact. Launched in 2008, the division takes a holistic, systems-oriented approach, working in areas from basic network services to architectures, data management abstractions, computer security, user interfaces, human factors, and domain-specific algorithms. The division specializes in highly collaborative user-driven research, in which we evaluate our work in the context of operational, high-impact domain science.

In earlier work, we developed grid computing infrastructures to support the creation and operation of "virtual organizations" as a foundation for collaboration and discovery. This work focused on understanding methods for sharing computing and storage infrastructure across distributed resource providers and collaborators. The resulting methods played a role in two Nobel prizes—e.g., all the data analysis for discovering the Higgs boson was performed on a global grid infrastructure, and the recent discovery of gravity waves took place on a data grid.

More recently, the division has focused on biomedical applications. Our current collaborations cover a broad range of applications—from basic science to clinical use cases spanning molecular biology, basic neuroscience, neuroimaging, stem cell research, and cranio-facial dysmorphia.

Our researchers work closely with ISI's highly regarded artificial intelligence, networking, and distributed systems experts, as well as with two of USC's nationally ranked Viterbi School of Engineering units: the Daniel J. Epstein Department of Industrial and Systems Engineering, and the Department of Computer Science. A major current initiative, the first direct study of living synapses in the intact brain, is an interdisciplinary effort with USC's Keck School of Medicine, Dornsife College of Letters, Arts and Sciences, and the Viterbi School of Engineering. In addition, the division plays a central role in four international consortiums, as well as one of the eight National Institute of Health funded centers of excellence for Big Data.

Intelligent Systems Division | *Director Yigal Arens*

The Intelligent Systems Division (ISD) comprises more than 100 faculty, research staff, USC graduate students, and short– and long–term visiting researchers. Most ISD researchers hold graduate degrees in computer science or related disciplines, and many also serve as research faculty in the USC Viterbi School of Engineering – mainly in the Department of Computer Science.

Members of the Intelligent Systems Division at their biannual retreat

ISD is one of the world's largest artificial intelligence (AI) groups. It is known especially for its work in natural language processing, machine translation, information integration and social networks. We also explore biomedical data integration and engineering, computational behavior, adaptive robotics,and video, image and multimedia analysis. We build working prototypes and partner with industry to create commercial applications. Below are ISD's primary research thrusts.

Natural language processing and machine translation, for which ISD is internationally renowned; this includes statistical machine translation, question answering, summarization, ontologies, information retrieval, poetry generation, text decipherment, and more.

Knowledge technologies, involving workflow creation and management for very large-scale scientific and other applications.

Information integration, using artificial intelligence and machine learning techniques to solve complex information integrations problems, with applications to areas ranging from describing artwork to discovering exploited children.

Biomedical data integration that provides a semantically consistent view of, and efficient query access to, distributed, heterogeneous biomedical data.

Biomedical knowledge engineering, and innovative methods of building biomedical informatics systems based on cutting-edge AI techniques.

Robotics, in particular modular, self-reconfiguring robots and control methods.

Social networks that decipher social web structure and dynamics from multiple perspectives.

Video, image and multimedia analysis, including document image processing and face recognition.

Internet & Networked Systems Division | *Director John Wroclawski*

ISI's Internet and Networked Systems Division has provided key impetus to the advance of world–changing technologies from the earliest days of the Internet to the present.

Today's group carries out a broad program of research spanning computer networking, cybersecurity, and interaction between the cyber and physical worlds. Key areas of study include:

A unique experimental methodologies research program—*"research on research"* in the large-scale complex networked systems space. This goal here is to improve the *research process itself* across our focus areas, providing researchers nationally and worldwide with powerful new tools for experimental cybersecurity and networked systems research.

Cyber physical systems research focused on the smart grid and similar large-scale problem domains. Building on the key observation that energy, transportation, and similar physical-world networks share many properties with cyber networks, our research seeks to apply cyber-network principles to the creation of robust, resilient, real-world networked critical infrastructures for modern society.
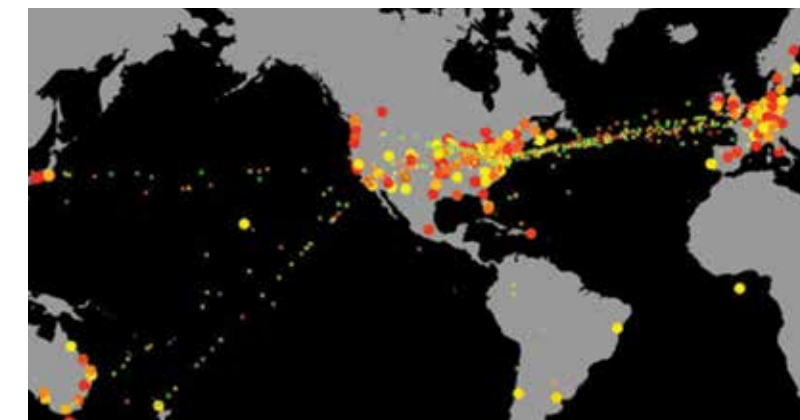
A sensing and sensor-net program focused on "self sustaining sensing"—that is, sensing networks for the physical world that operate autonomously, sustain themselves indefinitely, and function without human intervention. In applications ranging from oilfield monitoring to underwater environmental sensing, self-sustaining sensor nets transform both the capabilities and the economics of modern cyber-physical systems.

A networked systems cybersecurity portfolio centered on large-scale and system-level challenges. Focusing on current cybersecurity challenges facing the Internet, our research addresses topics ranging from distributed denial-of-service to usable security solutions for safe web browsing.

A world-recognized Internet measurement group whose work in this area seeks to improve our fundamental ability to understand the Internet's operation and evolution over time, as it continues to evolve into a core structure of modern society.

New, "clean-break" initiatives in areas ranging from optical networking to software systems integrity. Here, our overriding goal is to "try something different"—to look at established problems through new lights, with the aim of radically altering the capabilities or economics of the best-known solution.

A defining element of our research philosophy is the leverage of powerful cross-couplings between research itself and the tools for performing it. Beyond pure research, we frequently design and build cutting-edge research infrastructure—as catalyst and enabler for our own work, as a transition path for our results, and as a service to the larger community. DETERLab and AmLight, pictured here, are but two examples of this synergy.

ISI's DETERLab provides advanced experimentation capabilities for the cybersecurity research community as a national resource. Shown is DETERLab modeling a large-scale Internet distributed denial of service (DDoS) attack.

AMLIGHT 2015
Consortium

The aim of the AmLight Consortium's ExP project is to ensure that it meets the needs of the science communities it serves.

MOSIS Division | *Director Wes Hansford*

*For 35 years, custom IC designers have relied on the **MOSIS service at ISI** for an efficient, affordable way to prototype and volume–produce their devices. Since 1981, MOSIS has processed an average of more than seven IC designs per day.*

Many turn to MOSIS for our special expertise in providing multi-project wafers (MPWs) and related services that drive IC innovation. This "shared mask" model combines designs from multiple customers or diverse designs from a single company on one mask set. It's a practical prototyping channel that allows designers to debug and perform essential design adjustments before making a substantial strategic investment. Today, with mask costs soaring, more designers than ever are using MPWs to manufacture proven devices and prototype new designs on a single wafer.

Beyond MPWs, customers are increasingly choosing MOSIS as their resource partner for volume-production. From design spec interpretation through mask generation, device fabrication and onto assembly, MOSIS is their trusted expert interface to the semiconductor ecosystem. In addition to our commercial service, MOSIS is part of the following active research programs:

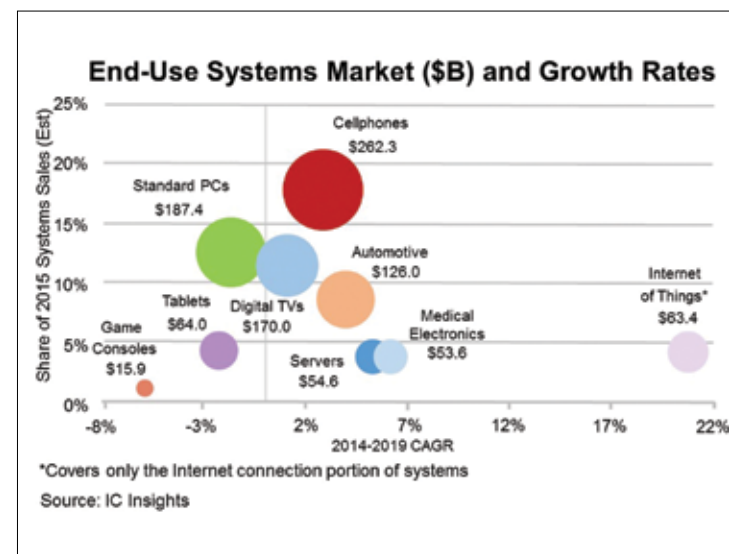**DARPA CRAFT (Circuit Realization At Faster Timescales)**— Developing a custom integrated circuit design flow and methodology that will: 1) sharply reduce the amount of effort required to design high-performance custom integrated circuits, 2) greatly facilitate porting of integrated circuit designs to secondary foundries and/or more advanced technology nodes, and 3) strongly increase reuse of integrated circuit elements. In support of CRAFT, MOSIS is organizing 14/16nm private MPW runs.

DARPA DAHI (Diverse Accessible Heterogeneous Integration)—Developing transistor-scale heterogeneous integration processes to intimately combine advanced compound semiconductor (CS) devices, as well as other emerging materials and devices, with high-density silicon complementary metal-oxide-semiconductor (CMOS) technology. The ultimate goal of DAHI is to establish a manufacturable, accessible foundry technology for the monolithic heterogeneous co-integration of diverse devices and complex silicon-enabled architectures on a common substrate platform. MOSIS is teamed with Northrop Grumman Corp., and is organizing private MPW runs (65/45nm).

IARPA TIC (Trusted Integrated Chips)—Developing and demonstrating split-manufacturing, with a new approach to chip fabrication where security and intellectual property protection can be assured. MOSIS is on the government team, and is organizing private MPW runs (130/65/28nm).
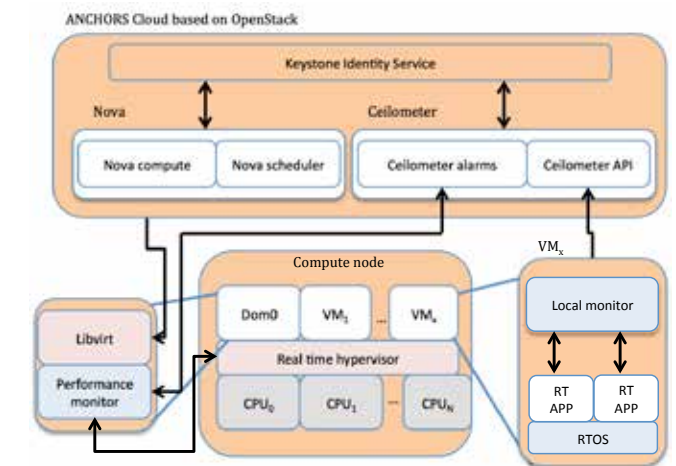

End-Use Systems Market ($B) and Growth Rates
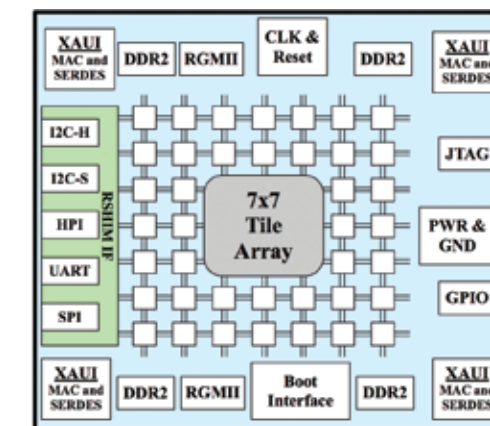
COMPUTATIONAL SYSTEMS AND TECHNOLOGY

ANCHORS: Autonomous Navy Cloud Supporting Heterogeneity & Optimized for Real-Time Systems

With few exceptions, today's cloud is predominantly homogeneous, supporting primarily commodity processor architectures and general-purpose workloads. This model is effective for many common applications, but is poorly suited for real-time workloads. ISI's *ANCHORS* project is developing a cloud and dynamic resource management framework for real-time mission-critical workloads based on the OpenStack cloud platform. The ANCHORS framework dynamically adjusts resource allocation to real-time workloads and hypervisors, adding and removing allocated processing resources at run time in order to increase server utilization while still meeting real-time quality-of-service requirements (e.g., deadlines). This is accomplished through the use of a runtime monitor within each real-time virtual machine and a performance monitor that adjusts resources dynamically. The ANCHORS project anticipates building on this work by improving resource management and extending the resource manager to address fault tolerance for real-time workloads.

ANCHORS is a component of the cloud computing group at ISI who are finding ways to integrate heterogeneous computing into dynamic public and private secure cloud platforms for data centers, high-performance computing, and embedded computing. ISI is one of the first organizations to develop support for heterogeneous computing in the OpenStack cloud environment, and is now developing support for real-time applications and the deployment of such technology to new application domains.

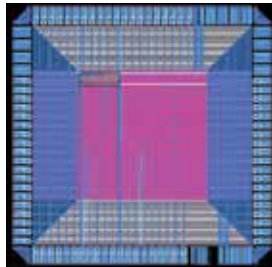*The ANCHORS cloud and real-time resource management framework*

Maestro: A Radiation-Hardened Space Processor

As Earth-orbiting satellites generate increasingly large datasets, new approaches to data ingest and processing must be considered. One approach is to increase the satellite-to-Earth downlink capacity, and process these datasets on the ground. A second approach, however, would allow the satellite to perform processing and/or preprocessing onboard. This would reduce downlink requirements, and may improve the timeliness of the data through autonomous satellite operation.

ISI's *Maestro* processor is a 49-core radiation-hardened-by-design (RHBD) processor for space missions. Maestro is derived from the Tilera TILE64 processor, inheriting a full complement of high-speed interfaces, a complete Linux distribution and a development toolchain. Our team of researchers is developing system software for the Maestro processor, including enhancements to the processor's compiler.

The Maestro Processor: Designed for space missions

Most visibly, USC/ISI is extending the Maestro toolchain with GCC compiler support, which will enable full OpenMP support and additional language front-ends. The switch to GCC has already resulted in improved performance for many applications. In addition, both OpenMP and the additional language front-ends will dramatically improve software portability to the Maestro processor.

COMPUTATIONAL SYSTEMS AND TECHNOLOGY
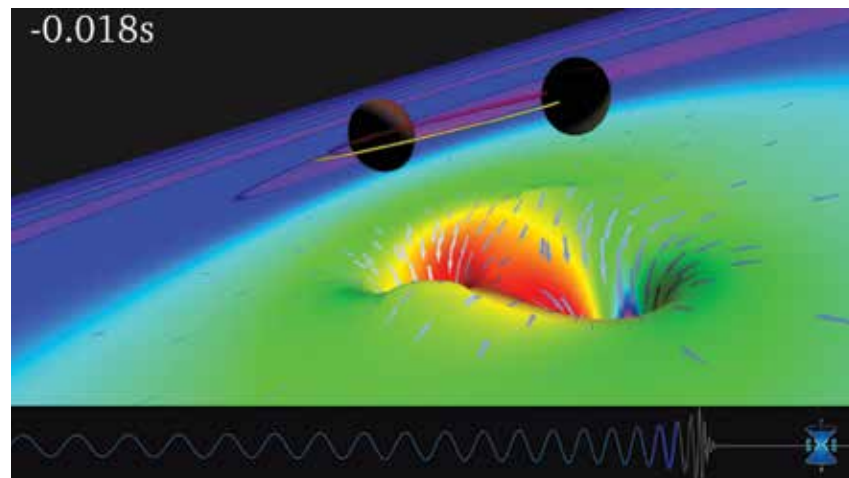

The RadMAC prototype chip

Radiation Mitigation Through Arithmetic Codes (RadMAC)

Sponsored by the Defense Threat Reduction Agency, the *Radiation Mitigation Through Arithmetic Codes (RadMAC)* project has been exploring the effectiveness of residual arithmetic codes in providing some measure of radiation mitigation in integrated circuits. In 2015, the RadMAC project designed, fabricated, and tested its first functional prototype chips, implemented in Samsung 28nm technology. The prototype chip depicted here was demonstrated to detect and correct errors as designed when exposed to radiation applied via a two-photon absorption laser source.
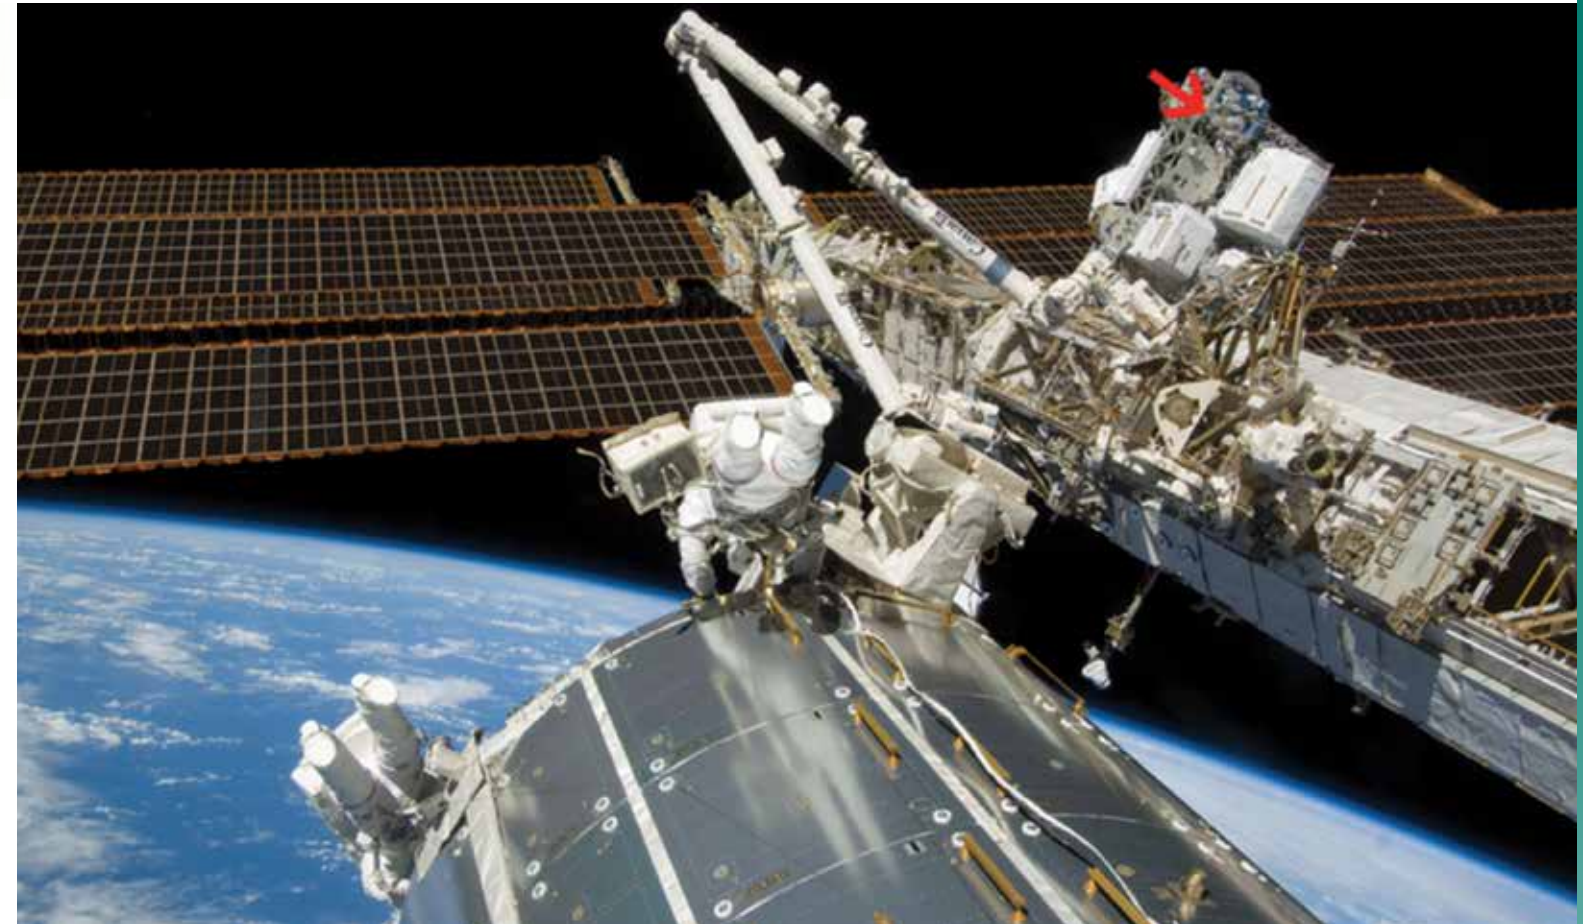
Pegasus: Powering the LIGO Gravitational Waves Detection

As early as 2001, the Pegasus team at ISI was working to create software that would accelerate scientific discovery by enabling researchers to analyze massive amounts of data. Early this year, physicists at the Laser Interferometer Gravitational-Wave Observatory (LIGO) announced a long-awaited scientific discovery: the detection of gravitational waves as predicted by Einstein's Theory of Relativity. *Pegasus* has been involved with LIGO since its early stages, when both groups became part of a National Science Foundation-funded project to explore the use of workflow technologies to run large-scale science pipelines on computational grids. The Pegasus team enabled the analysis of data gathered from the two LIGO detectors. Initially, pipeline analyses were managed by the Pegasus Workflow Management System on the LIGO Data Grid. LIGO then extended their computations to nationwide cyber-infrastructures, such as NSF's Open Science Grid (OSG) and Extreme Science and Engineering Discovery Environment (XSEDE). Pegasus contributed to this expansion by managing cross-site data transfers over wide-area networks, and large-scale computations in a reliable, scalable, and efficient manner.

One of the main pipelines (PyCBC) measures the statistical significance of data needed for discovery. This pipeline is composed of hundreds of thousands of computational jobs, and operates over terabytes of data. For gravitational waves detection, the pipeline used computational resources from Syracuse University and the Albert Einstein Institute-Hannover. The benefits of using Pegasus to run LIGO analyses are many: 1) an expanded computing horizon—Pegasus enabled LIGO pipelines to run across external cyber infrastructures beyond the LIGO collaboration, and 2) smart data management—Pegasus' optimization strategies enabled automated data discovery, automated data cleanup to reduce peak storage requirements, data reuse to reduce the amount of storage required for the computation, and support checkpointing.

0.018 seconds before the black holes collide (Image credit: SXS/LIGO)

The red arrow points to SpaceCube 2.0 hardware running a RHBSW flight experiment on ISS.

Autonomous On-Board Processing for Sensor Systems STP-H4 Flight Experiment
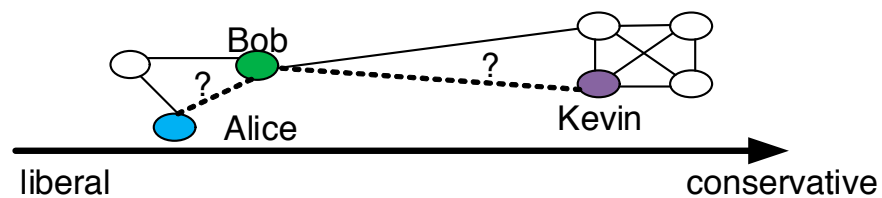
Under a series of NASA research grants, ISI's Reconfigurable Computing Group developed novel radiation mitigation techniques that enable software to detect and correct radiation-induced errors in commercially available processors. This approach allows these processors to operate in the space environment, as opposed to using traditional radiation-hardened processors which are typically two to four generations behind the state-of-the-art and 20x more expensive. Under our latest project, *Autonomous On-board Processing for Sensor Systems* STP-H4 Flight Experiment, radiation hardening by software (RHBSW) techniques, developed for embedded PowerPCs within Xilinx Virtex-5 FPGA devices, were demonstrated on board the International Space Station (ISS). Our RHBSW techniques leverage concepts such as assertions, checkpointing and rollback, commonly used in the high-performance computing industry, and adapts them for use in remote sensing embedded systems. These techniques have extremely low overhead and thereby enable a 3.3x gain in processing performance compared to the equivalent traditional radiation-hardened processor.

Previous research performed laboratory-and laser-induced failure injection campaigns to achieve a technology readiness level (TRL) of 6. In-flight experimentation is a critical step to achieving higher TRLs where technology can be transitioned to science missions. Under this effort, we ported the techniques to NASA GSFC's SpaceCube 2.0 hardware, which already resides on the International Space Station, developed the flight software with NASA, and remotely uploaded the new experiment, which ran for 16 days before the platform was decommissioned. The RHBSW executed on two PowerPCs embedded within the Virtex5 FPGA devices. During the experiments, we collected 19,400 checkpoints, processed 253,482 status "heartbeats," and incurred zero faults.
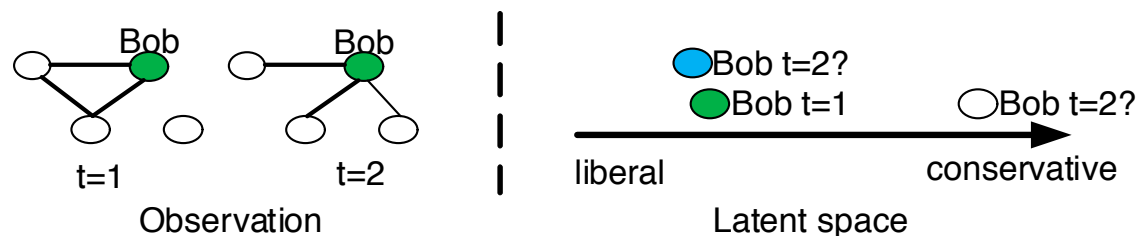
INTELLIGENT SYSTEMS

Scalable Link Prediction in Dynamic Networks via Non-Negative Matrix Factorization

This project seeks to understand and characterize the processes driving social interactions, which are key to the fundamental problems in social network research. A particular instance of this problem, known as link prediction, has recently attracted considerable attention in various research communities. In addition to being of purely academic interest, *link prediction* has many important commercial applications, e.g., recommending friends in an online social network such as Facebook and suggesting potential hires in a professional network such as LinkedIn.

An example of observed interactions among Alice, Bob and Kevin, and their positions in a simplified one-dimension latent space. Alice is very liberal, Bob is biased towards being liberal, and Kevin is very conservative.

An example of temporal latent positions of user Bob. With the temporal smoothness assumption, Bob is more likely to be in the blue position than the white position in time t=2.

Automatically Extracting Pedagogically Valuable Reading Lists for Technical Corpora (TechKnAcq)

The *TechKnAcq Project* is concerned with developing scalable methods for preparing pedagogically valuable reading lists from technical corpora in response to user queries. This goes beyond simple search technology since the returned list of documents must be appropriately sorted to reflect conceptual dependencies in the raw text of a given corpus. For example, in a reading list generated for students working on machine learning, prerequisite reading materials should be included in the reading list *before* the core documents describing the elements that reflect the closest hits to the query. We are also developing methods for scoring material based on its pedagogical value to students.

Large-Scale Face Recognition in the Wild @ ISI

Established in the fall of 2013, the Computer Vision Group at ISI has quickly grown into a vibrant team that is advancing the state of the art in areas such as text extraction from images and videos, optical character recognition, and face recognition. Sponsored by IARPA's Janus program, the face research effort (*GLAIVE*, for *Graphics-based Learning Approach Integrated with Vision Elements*) is a close collaboration between ISI's Computer Vision Group and the IRIS Lab at USC's University Park campus.

The purpose of IARPA's Janus program is to develop face recognition systems capable of handling images and videos observed "in the wild," i.e., images and videos that are not acquired under controlled conditions and are for the specific purpose of facilitating face recognition. As a result, the developed algorithms need to handle a wide range of variations in attributes, such as face pose, age, and expression, as well as varying image qualities and illumination conditions. Furthermore, algorithms and systems developed under Janus must be capable of handling very large numbers of database subjects, and enrolling new subjects into the reference database without retraining the underlying machine learning algorithms. Subjects' data could come from various sources, such as images, videos, infrared and sketches, and each subject is represented by one or more facial images, as shown in this accompanying collection of images.

Sample images from IARPA Janus dataset

Our approach combines algorithms from computer graphics, deep learning, and computer vision in developing a state-of-the-art face recognition system. Key aspects of our approach include graphics techniques that use 3D face models to generate unseen face views (or poses) from a limited number of actual "seen" views, convolutional neural networks, transfer learning to improve in-domain performance by leveraging out-of-domain data, locality-sensitive hashing to improve matching speed, and a novel representation called SQSH, Sparse Quantized Scalable Hyper-volume, which facilitates a hierarchical fusion strategy. The ISI face recognition system has been evaluated on the recently released NIST IJB-A dataset, and it demonstrated state-of-the-art performance.

Communication with Computers

Sponsored by DARPA's *Communication with Computers Program*, this research effort contends that as machines (computers and robots) become increasingly intelligent, we need to learn how to communicate with them using regular English; further, they need to be able to explain what they are up to, again using regular English. We will be stress-testing our research assumptions and algorithms in the context of two applications: 1) In collaborative writing, we will develop models that understand why a sequence of sentences makes up a coherent text, and can generate enticing narratives. We will apply these models to understand/generate stories and poetry. 2) In human-robot communication, we will develop algorithms that enable machines to understand simple and complex human instructions. This requires the meaning of words to be grounded in the physical worlds in which robots operate, and algorithms that enable robots to explain what they do and plan to do in plain English.

INTELLIGENT SYSTEMS

Efficient Estimation of Mutual Information for Strongly Dependent Variables

ISI researchers have demonstrated that a popular class of non-parametric mutual information (MI) estimators based on k-nearest-neighbor graphs requires a number of samples that scale exponentially with the true MI. Consequently, accurate estimation of MI between two strongly dependent variables is possible only for prohibitively large sample sizes. This important yet overlooked shortcoming of the existing estimators is due to their implicit reliance on local uniformity of the underlying joint distribution. We introduce a new estimator that is robust to local non-uniformity, works well with limited data, and is able to capture relationship strengths over many orders of magnitude. We demonstrate the superior performance of the proposed estimator on both synthetic and real-world data.

Machine Reading of Cancer Experimental Observations and Methods

As a performer in DARPA's *Big Mechanisms* program, researchers from the Intelligent Systems Division are developing methods for machine reading of experimental observations from research papers. This work is motivated by the idea that within biomedical research, measurements are essentially objectively accurate, and interpretations based on those measurements are generally more subjective. In other words, if the same experiments are performed, it is likely that the same measurements will be obtained; however, it is less probable that two people will make the same interpretations of those measurements. In general, researchers using information extraction (IE) on biomedical papers focus on "main punch lines" from papers rather than the nuts and bolts of experimental outcomes. This makes sense to NLP researchers since IE methods are somewhat unreliable, and the focus is on extracting only information that is of the most high value. We argue that the primary findings of papers contain information that is likely to be the most contentious, and this also does not reflect the way that biologists read papers (where the typical focus is on experimental details rather than high-level arguments).

The focus for Big Mechanisms is on automatically reading the content of papers pertaining to cancer pathways; other members of the program focus on reading and synthesizing the main findings of papers ("protein A binds protein B," or "gene X activates protein Y"). Our work focuses on extracting the underlying evidence that supports such claims. So far, we have developed and tested IE technology that can classify clauses in the results section of papers based on the type of discourse structure they invoke.
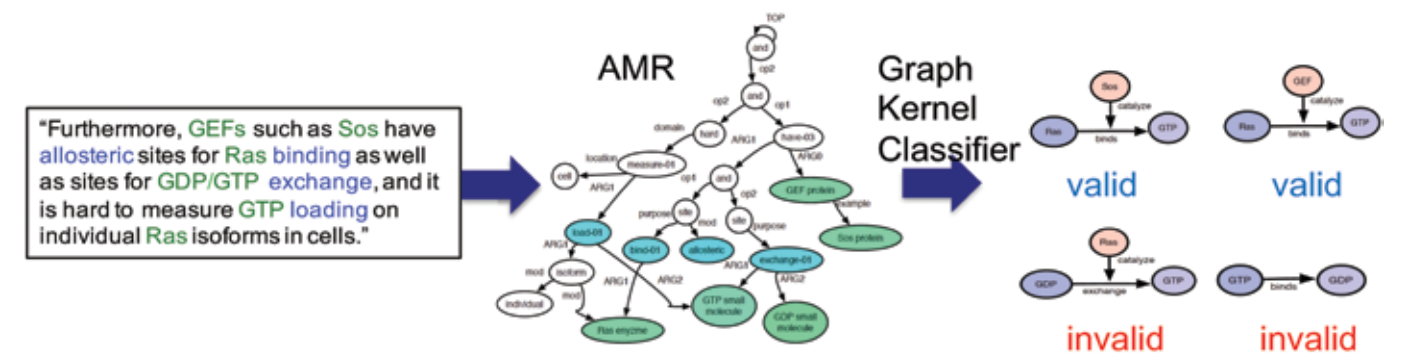
Knowledge Engineering from Experimental Design (KEfED)

A component of the work targeted for the *Big Mechanisms* program uses a method of modeling the semantic structure of scientific experiments called "knowledge engineering from experimental design" (KEfED). We are using this approach to develop models of the data structure used in systems biology studies of intracellular signaling pathways in cancer cells. This work directly builds on machine reading efforts and provides a more concrete, structured backbone for that work. In particular, our ISI research group is performing detailed modeling of experimental protocols in cancer experiments to construct a generative modeling architecture for experimental types. This will support our information extraction work and will also contribute to the development of knowledge representation and reasoning technology for molecular biology data.

Extracting Biomolecular Interactions Using Semantic Parsing of Biomedical Text

Biomedical knowledge extraction suffers from several shortcomings. First, most current methods rely on shallow analysis techniques that severely limit their scope — e.g., focusing on whether there is an interaction between a pair of proteins while ignoring the interaction types. Second, these methods also focus on single-sentence extraction, which makes them very susceptible to noise. And third, owing to the huge diversity of research topics in biomedical literature and the high cost of data annotation, there is often a significant mismatch between training and testing corpora; this reflects poorly on the generalization capabilities of existing methods.

Our research team has developed a novel algorithm for extracting biomolecular interactions from unstructured text that addresses the above challenges. Contrary to previous work, the extraction task considered here is less restricted and spans a much more diverse corpus of biomedical articles. These more realistic settings present some important technical problems for which we provide explicit solutions: 1) a graph-kernel based algorithm for extracting biomolecular interactions from abstract meaning representation (AMR); 2) multi-sentence generalization of the algorithm by defining graph distribution kernels for performing document-level extraction; 3) a hybrid extraction method that uses both AMRs and syntactic parses given by Stanford dependency graphs.

A sentence is parsed into abstract meaning representation; then a graph kernel is applied to extract valid interactions.

Pediatric Research using Integrated Sensor Monitoring Systems (PRISMS): Data and Software Coordination and Integration Center (DSCIC)

The Pediatric Research using Integrated Sensor Monitoring Systems (PRISMS) Data and Software Coordination and Integration Center, funded under the National Institute of Biomedical Imaging and Bioengineering, is a four-year effort to develop an architecture for the integration and analysis of environmental, physiological and behavioral data produced by the PRISMS sensor-based health monitoring systems (and other systems). The primary domain is the epidemiological studies of asthma, and eventually other chronic diseases in the pediatric population.

The goals of the Center are to: 1) integrate the multiple data sources produced by the PRISMS program into a well-described, semantically consistent dataset; 2) link PRISMS data to outside knowledge bases and resources; 3) provide large-scale analytics over the integrated data; and 4) build innovative access and query tools to allow for seamless exploration of the data by a broad range of biomedical researchers.

The Center's research team includes faculty and researchers from USC's Information Sciences Institute, the Keck School of Medicine, and the Department of Computer Science.

INTELLIGENT SYSTEMS

OntoSoft: Software Metadata

Scientific software captures important knowledge and should be shared and reused. Although there are many popular code repositories, they do not contain descriptions that help scientists find and reuse software. In addition, there is still a significant amount of scientific software that is never shared, in part because code repositories require more programming skills than scientists have.

The *OntoSoft* Project is developing a software registry where software is described with metadata that is designed with scientists in mind. The metadata captured by OntoSoft is organized in an ontology of six major categories. This organization allows the scientists who contribute software to understand why the metadata is requested, and to allow the scientists looking to reuse the software to find metadata that is relevant to them.

OntoSoft is being used in the NSF EarthCube program for geosciences. It is deployed as a federated registry where different communities run their sites, but they are all connected so users can conduct searches across all sites. OntoSoft also allows a software author to give permission to other users to extend the metadata of his software based on their experiences with its use.

Geoscience Papers of the Future (GPF) Initiative

In the next 5 to 10 years we expect that scientists will use radically new tools to develop research papers. These tools will document and publish all the associated digital objects (data, software, physical samples, etc.) that form the basis for the paper, together with detailed provenance and workflow for all new research results. This evolution in research publications will substantially improve science communication, promote a fair basis for crediting science contributions, and offer a transparent way for other scientists to evaluate and reproduce the research. It will also make research products machine-readable, enabling use by intelligent systems.

To help make this vision a reality, the *Geoscience Papers of the Future (GPF)* Initiative is a project led by ISI researchers to promote the publication of geosciences papers together with the associated digital products of the research. The GPF initiative promotes best practices for:

Open science, ensuring that the research methods and products are publicly available in a manner that allows others to inspect, reproduce, and reuse any aspects of the research.
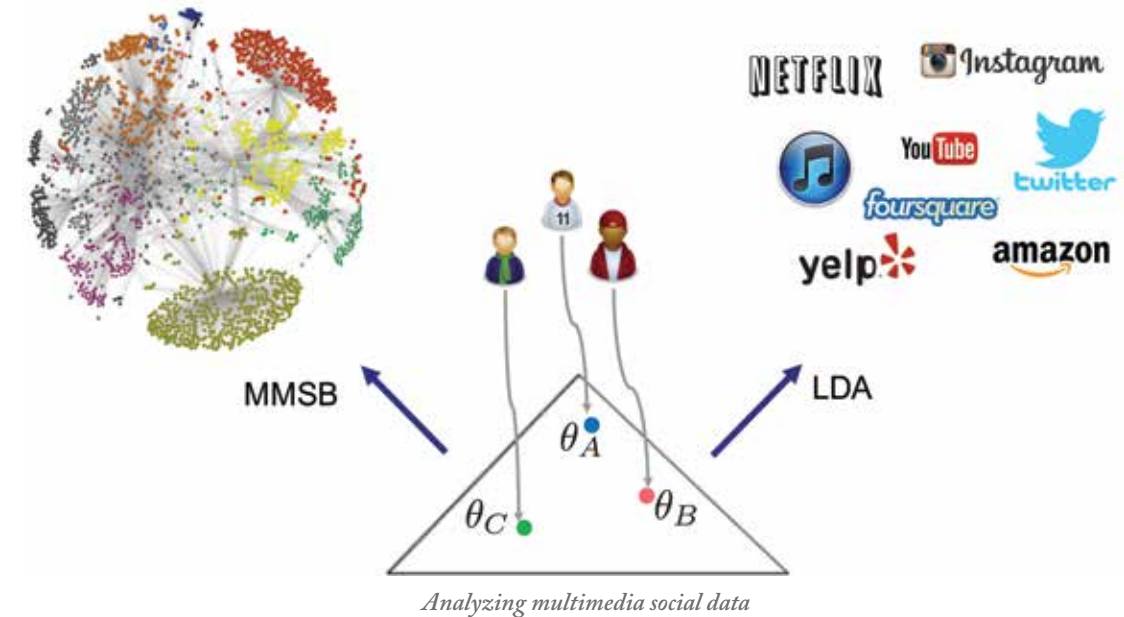
Data and software stewardship, emphasizing that research products (data, software, and methods) should be managed, documented, and archived in sustainable repositories as important contributions to science.

Digital scholarship, recognizing that scientific publications should include citations and unique identifiers (DOIs) for datasets, software, physical samples, and other research products in order to preserve their identity and enable authors to receive credit for them.

The GPF Initiative has developed training materials to teach best practices and conducted training sessions at major geosciences events that included more than 250 participants.

Latent Space Models for Multimodal Social Data

With the emergence of social networking services, researchers enjoy the increasing availability of large-scale heterogenous datasets capturing online user interactions and behaviors. Traditional analysis of techno-social systems data has focused mainly on either the dynamics of social interactions, or the attributes and behaviors of the users. However, empirical evidence suggests that the two dimensions affect one another, and therefore should be jointly modeled and analyzed in a multimodal framework. The benefits of such an approach include the ability to build better predictive models, leveraging social network information as well as user behavioral signals. Developed by an Intelligent Systems Division team of researchers, the *Constrained Latent Space Model (CLSM)* is a generalized framework that combines mixed membership stochastic blockmodels (MMSB) and latent Dirich-let allocation (LDA) incorporating a constraint that forces the latent space to concurrently describe the multiple data modalities, ultimately making it feasible to analyze massive social datasets.

Analyzing multimedia social data

Correlation Explanation for Multi-Model Brain Data Analysis

Total Correlation Explanation, or CorEx for short, is a new information-theoretic method developed by ISI's researchers that can be used for discovering relationships in complex systems. CorEx has been shown to recover meaningful structure in diverse systems that include human behavior and language, and biomedical domains. An ongoing collaboration with the Imaging Genetics Center has recently leveraged CorEx to discover novel biomarkers predicting Alzheimer's disease. Another effort has discovered new signals in gene expression data that predict long-term survival for ovarian cancer patients.

INTELLIGENT SYSTEMS

"Majority Illusion" in Social Networks

An individual's attitudes and opinions are shaped by his/her perceptions of the attitudes and opinions of others. As a result, these can spread through the population as people perceive their friends adopting a particular attitude or behavior, and then adopting it themselves. A group of ISI researches are studying this kind of behavior in social networks.
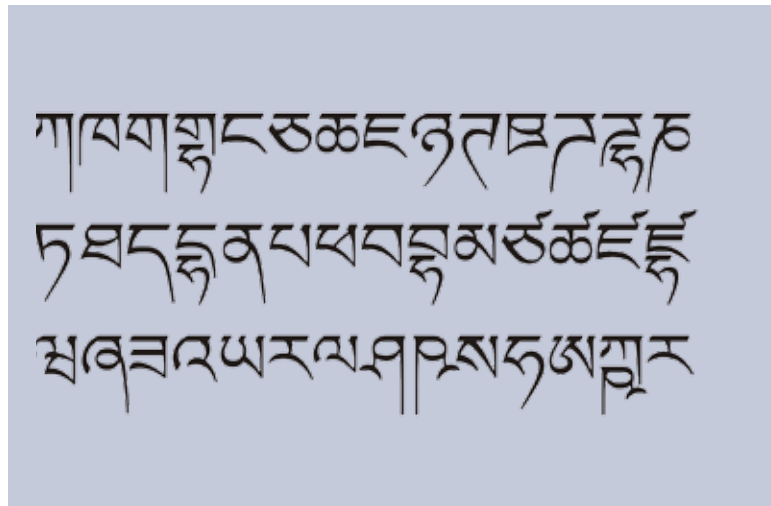
The network structure itself can systematically bias social perceptions people make about their peers. One of the better known of these biases is the *friendship paradox*, which states that, on average, most people have fewer friends than their friends have. Recently, our researchers identified a novel variation of the friendship paradox, called the *majority illusion*. The paradox creates conditions wherein many people in a social network observe that a majority of their friends have some trait, even if it is globally rare.

As a simple illustration of the "majority illusion," consider these two networks. The networks are identical except that different nodes in the two networks are colored. Despite this apparently small difference, the two networks are profoundly different. In the first network, every white node will examine its network neighbors (i.e., friends) to observe that "at least half of my friends are red," while in the second network no white node will make this observation. Thus, even though only three of the 14 nodes are red, it appears to all white nodes in the first network that most of their friends are red.

ELISA: Natural Language Processing

Natural language processing (NLP) capabilities are currently limited to a handful of the world's 4000+ languages. ISI's new *ELISA Project* anticipates a natural or man-made disaster somewhere in the world, followed by a torrent of tweets and news in the local language. Our goal is to develop technology that will quickly field NLP tools for that local language. This means being able to translate text, extract important entities and their relationships, and detect sentiment/emotion in both text and speech, with the overall aim of providing direction for humanitarian assistance. We are carrying out this project under the auspices of the DARPA *LORELEI Program*, together with collaborators at the University of Pennsylvania, the University of Notre Dame, Next Century Corp., Rensselaer Polytechnic Institute, Brno University of Technology, and the International Computer Science Institute (ICSI).
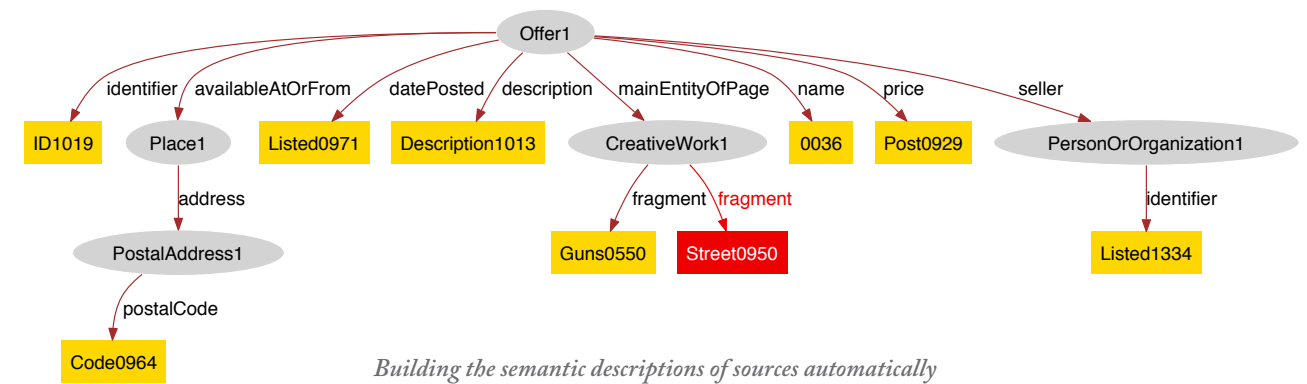

Tibetan language—just one of the world's 4000+ languages

Karma: Automatically Building Semantic Descriptions of Data Sources

Over the past decade, ISI researchers have been working on a simple and fast way to integrate data from heterogeneous sources. Towards that goal, we developed an integration tool called *Karma*, which can rapidly integrate data from a variety of sources. Karma's unique capability is that it learns semantic descriptions of sources. These descriptions define a mapping from the contents of a source (e.g., all the columns in a spreadsheet) to a domain ontology provided by a user. Karma uses these descriptions to align, integrate, and query data across sources.

This past year, Karma researchers focused on building the semantic descriptions of sources automatically. They developed scalable and accurate techniques for identifying the mapping of an attribute to the ontology and formed an approach that combines numerous features using a machine learning technique called *random forests*. This technique learns a decision tree based on a variety of features and then classifies the attribute based on values of the features. Then, since our researchers also needed to determine how attributes are related to each other, they developed an approach that learns these mappings from information already published on the Web. The idea behind this is to mine the patterns in already-existing data and then use those patterns to predict the relationships in new data sources.

Our researchers are now applying the automated techniques to one particularly compelling problem—that is, the automatic extraction and modeling of data from sources that are advertising potentially illegal arms sales. The idea is to start with a website that is selling guns, collect all of the pages from that site, automatically extract and group the data from those pages, and then apply the Karma techniques to identify the fields on the web pages and the relationships between those fields. We evaluated the technique on 10 websites, and the results are quite promising.

Building the semantic descriptions of sources automatically

AI in Education (Short-Answer Assessment)

Computers can easily assess the correctness of a multiple-choice answer via a simple lookup. However, when answers are written in regular English and are several sentences long, automatically assessing how good an answer is becomes difficult. Under sponsorship of Educational Testing Services, ISI is developing novel algorithms for automatically scoring answers to short questions. We do this by applying techniques that have proven useful in the context of automatically scoring the proficiency of human translators.

INTELLIGENT SYSTEMS

Teaching Data Science to Non-Programmers

Big data analytics has emerged as a widely desirable skill in many areas. Although courses are now available on a variety of aspects of big data, there is a lack of broad and accessible courses for non-programmers that enable them to learn about big data analytics in practice.

With a grant from the National Science Foundation, ISI researchers have developed course materials to train non-programmers in important data science topics: machine learning, parallel computing, visualization, ontologies and semantic metadata, and data stewardship. Students learn to analyze text data, multimedia data, geospatial data, timeseries data, and network data. Analytic methods are captured in semantic workflows, which ensure valid use of the methods and provide an easy-to-use visual programming environment where students can practice within predefined lesson units. Students taking this course are in political science, education, communications, and business. It is being taught at USC in Spring 2016 as part of the Informatics Program, and will become a requirement for all new interdisciplinary degrees in Informatics at USC.
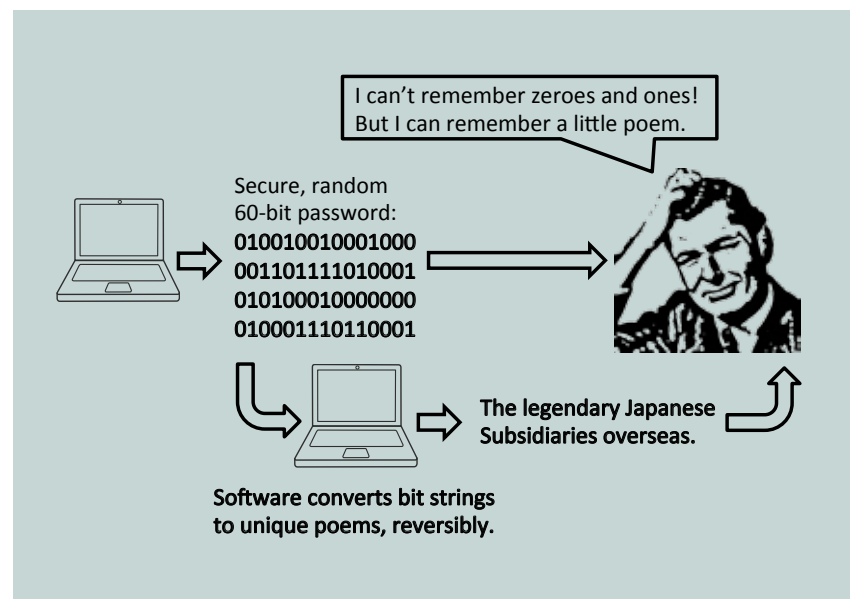
BigDataU: The NIH Big Data to Knowledge (BD2K) Training Coordination Center

The *BD2K Program* is the flagship Data Science initiative of the National Institute of Health, which has funded about a dozen major centers exploring big data management across different biomedical domains, ranging from neuroimaging, to genetics, to mobile health. An important component of this effort is to train biomedical researchers in data science.

In addition to the two existing USC BD2K Centers: the Big Data for Discovery Science Center and the ENIGMA Center for Worldwide Medicine, Imaging, and Genomics, NIH has awarded USC the *BD2K Training Coordination Center*. The goal of the BD2K Training Coordination Center is to develop an online education platform that will automatically discover, access, curate, cite, and personalize training materials in biomedical big data. The BD2K research team includes faculty from USC's Information Sciences Institute, Keck School of Medicine, School of Cinematic Arts, and Georgetown University.

Machine-Produced Poetry Passwords

ISI's Natural Language Processing Group recently discovered an exciting new application for their work. These researchers have created a software program called *Machine-Produced Poetry* Passwords that generates secure, memorable passwords in the form of short English-language poems, called couplets. The system works by first creating an initial, random 60-bit password. Because the password is random, it is secure against brute-force guessing in a way that human-invented passwords are not. However, since a long, random string of ones and zeroes is too difficult for any user to memorize, the group developed natural-language software that maps any such string onto a unique two-line poem. The poem is much more memorable, and because the mapping is reversible, the system maintains its secure properties.
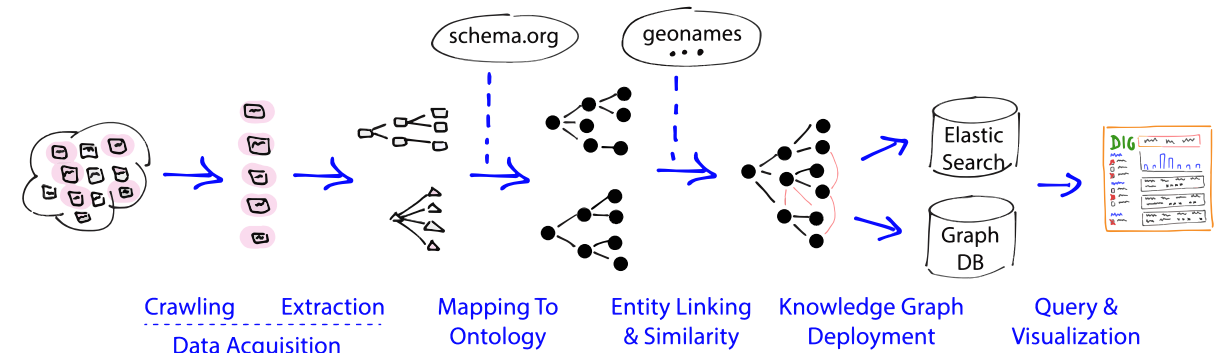

Creating secure, memorable passwords

Domain-Specific Insight Graphs

The *Domain-Specific Insight Graphs (DIG) Project* is a collaborative effort between USC/ISI, Columbia University, InferlinkCorp., and Next Century Corp., as well as multiple partners in the DARPA MEMEX program. The goal of the DIG project is to reorganize the web so that it can be used for investigative search. Conventional search engines such as Google and Bing leverage the links in pages to find the most popular and influential pages given a set of keywords. Investigative search is about finding information to solve a complex problem—such as finding victims of human trafficking, illegal firearm sales or counterfeit computer chips. These searches often involve finding outliers and implicit links among pages such as patterns of mentions of a phone number in multiple web pages.

To address these problems, DIG is developing techniques to 1) *identify* websites containing relevant information; 2) *develop* crawling technologies to download a corpus of relevant pages (potentially hundreds of millions of pages); 3) *extract* information from these pages using methods which automatically identify a common structure among pages that enables extraction of data; 4) *integrate* the information extracted using a common ontology; 5) *use entity* resolution to identify records that refer to the same entity in the real world using a combination of highly discriminative features such as names, and weak features such as photos, writing style; 6) *use NoSQL* databases to store knowledge graphs to enable sophisticated query and scalable analytic algorithms; 7) *use* image analysis to identify and categorize visually similar images; 8) *develop* user interfaces to enable investigators to easily query the resulting knowledge graph to solve complex problems.

Using DIG, ISI's researchers have created domain-specific knowledge graphs and user interfaces for a variety of domains, including illegal firearms trafficking, identification of patent trolls, analysis of research trends in material science and autonomy, and human trafficking.

THE DIG Project: Reorganizing the web for investigative search

Big Mechanisms for Bio-Molecular Modeling

Sponsored by DARPA's *Big Mechanisms Program*, this project has made significant progress towards automatically assembling complex bio-molecular models of cancer, starting from a trove of scientific papers. Our progress was driven by three catalysts: 1) the creation of a large corpus of abstract meaning representations (AMR) of texts in the bio domain which make explicit the semantics of bio-molecular interactions; 2) the development of novel, state-of-the-art AMR parsers that automatically infer bio-AMR representations for unseen texts; 3) the development of interaction extraction and assembly algorithms that create complex models of molecular interactions that computers can reason with. During 2015, we built a solid technology platform that now enables us to investigate what novel drug cocktails may provide adequate cancer treatment, as well as offer a deeper understanding of cancer pathways.

INTERNET AND NETWORKED SYSTEMS

CiSoft Pervasive Thermoelectric Generation (PTG) and CiSoft Sensor Network (SEN)

The *CiSoft Pervasive Thermoelectric Generation (PTG) Research Project* came to a conclusion in 2015 after five years of innovative research in efficient harvesting and conditioning of ambient thermal energy into usable power for industrial standard wireless SCADA devices. The CISoft PTG research team deployed working systems in the oil fields for the San Joaquin Valley Business Unit (SJVBU) of Chevron in Bakersfield, California. Currently, there is an on-going discussion regarding commercialization of the technology through a newly established commercialization entity, CiSoft Solutions, and possible partnerships with existing companies such as Perpetua, Emerson, and GE.

Deployed energy harvesting sensors for cold oil blockage detection research (CiSoft SEN)

Given this highly efficient and inexpensive technology, ISI researchers are continuing to work with Chevron's SJVBU through CiSoft to transition research activities to a new project—the CiSoft Sensor Network (SEN). We designed and built low-power and intelligent wireless sensors for oil field applications; this technology is intended to be the platform for research efforts in self-sustaining cyber-physical systems. We've instrumented two of these sensors on a production pipeline to continuously collect high-quality vibration/acoustic data without experiencing any problems. The sensor nodes harvested ambient energy from the environment to enable their instrumentation in remote locations where grid power was not available. Our current research goal is to analyze this data using a novel biomimetic audio processing method to detect and locate cold oil buildup within the pipelines before it causes costly pipeline blockage. Additionally, other oil field applications for this approach are being examined.

Improving DNS Privacy and Security: From Research in 2014 to Standards in 2016

The first step in retrieving a web page or sending an e-mail message is to *ask for the name* — the "isi.edu" in http://www.isi.edu and example@isi.edu. Today these requests are sent in the clear and replies are limited by packet size. These constraints undercut privacy and security by not using or limiting a choice of cryptographic algorithms.
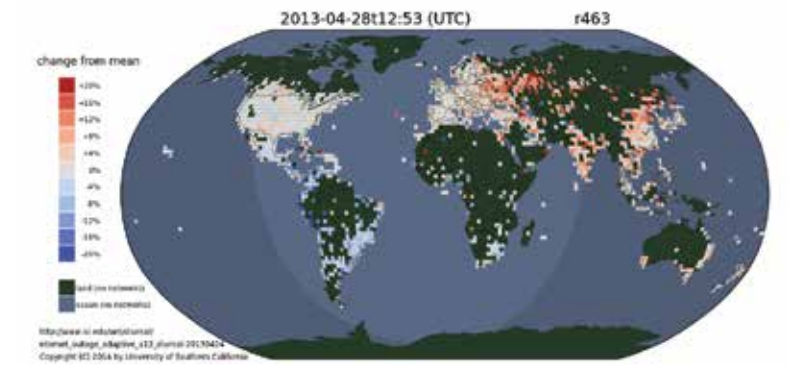
We propose to protect DNS privacy using TLS, the same protocol used in secure web connections (HTTPS). We have conducted experiments, simulation, and modeling to show that TCP and TLS have only modest performance cost (about 20% slower than UDP) when they are used in the right way. Our work helped disprove the myth that TCP and TLS are too slow for use with DNS, and resulted in an Internet standards document (RFC) to appear in early 2016. As our approach becomes a standard Internet protocol, we expect to see it increasingly used for DNS looks, particularly for people who use services like Google Public DNS.

Studying Internet Use: Ongoing Work

Of course people use computers; but not *all* the time — people sleep. *Computers sleep*, *too* — laptops are often suspended at night, and some computer labs are turned off. Here at ISI, researchers have studied Internet address use for years, but can we see the Internet sleep?

We conducted a study, using data collected in our evaluation of Internet outages. We found that *yes, parts of the Internet sleep*. As shown in our figure, there are more active addresses (red) during the day and at sunset, and fewer (blue) in the early morning darkness and just after sunrise. Other parts of the Internet are much more stable (always white).

We showed that it is possible to detect *short-term* variations in Internet use from a centralized, active prober. Our results prove the benefits of U.S. FCC policies to promote an always-on Internet as part of their broadband goals. We also observed that in each country, *greater sleep correlates with lower economic activity and electrical usage*; these observations can help us understand our world.

A sleeping Internet: blue areas are fewer addresses than normal, red are more.

Methodologies for the Study of Next-Generation Cyber Physical Systems

This research aims to fundamentally advance the state-of-the-art in computational and experimental methods for the study of large-scale, complex cyber-physical systems, and then to apply these methods to address critical infrastructure challenges across domains such as energy systems, smart cities, transportation, natural gas, and water supplies.

Recent technology advances in areas such as smart and ubiquitous sensors, communications, and distributed controls are rapidly driving the scale and complexity of interconnected cyber-physical systems and networks to new levels. Building on strong ties to the computer networking community, which has already confronted similar issues of extreme complexity, decentralized administration, and rapid, near-exponential growth, ISI researchers are leading the development of new modeling, experimentation, and analysis strategies for the rapidly advancing world of large-scale cyber-physical systems. Our research synthesizes ideas and methods from disciplines that include statistics, information theory, systems, optimization, economics, policy, network science, and control theory.

Our core goal is to create new methodologies and tools that will provide clear, correct, and actionable insights into the design and operation of these increasingly central components of modern society.

INTERNET AND NETWORKED SYSTEMS

DETER: Experimental Methodologies for a Secure Cyber Future

ISI is home to the *DETER Project*, an internationally recognized research effort sponsored by DHS, NSF and DARPA, whose goal to develop next-generation *methodologies, tools*, and *infrastructures* for experimental cybersecurity research. Rather than focusing on individual results, DETER seeks to advance state-of-the-art cybersecurity R&D *overall*, with particular application to complex, distributed, multi-organizational, networked cyber systems.

Central to this activity, the project developed and operates the DETERLab, a unique systems modeling and emulation environment established to support experimental cybersecurity research. The DETERLab serves as a national resource for the cybersecurity research, development, and education communities, and as an immediate technology transfer-to-practice vehicle for new research results.
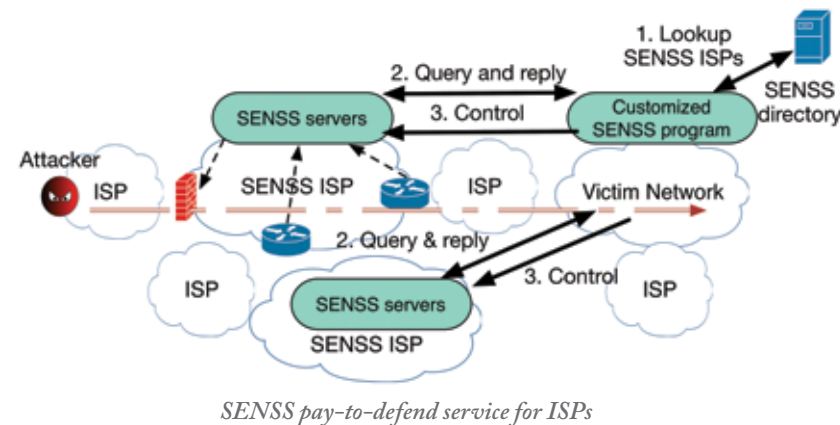
A recent key research advancement from this program is a wide-ranging collection of new, innovative capabilities for the modeling of *large-scale* and *highly complex* systems within an experimental research scenario. The DETER project has provided user communities with new capabilities that offer a unique approach to scale through multi-resolution virtualization, and a new class of unique, agent-based tools for managing extreme-scale experiment design, execution and analysis. These capabilities allow DETER researchers to model and study cyber-systems with hundreds of thousands, and soon, millions of elements.

Creating effective, hands-on strategies and approaches for students to study modern cybersecurity concepts and techniques is a challenge of worldwide interest. The DETER project has developed and made freely available a set of specialized support capabilities and teaching exercises for use in undergraduate and graduate cybersecurity classes. These advanced capabilities give faculty and students direct access to new, project-based forms of instruction, such as classroom "capture the flag exercises" that combine effective, compelling pedagogy with broad applicability and simple administration. DETER educational support is currently being used by over 8,000 students in institutions ranging from top-tier universities to small community colleges serving underrepresented populations.

Network Security and SENSS

With an increased reliance on the Internet for so many tasks in our daily lives and in business and critical infrastructure, network security is a topic of increasing importance. At ISI we work on understanding threats and designing effective defenses in network security: *passwords*—how people reason about security risk vs. memorability trade-off, and the design of new, more memorable authentication mechanisms; *distributed denial-of-service*—the design of deployable defenses

SENSS pay–to–defend service for ISPs

against high-volume attacks and application-level floods; *malware*—transparency mechanisms for VM, debuggers and testbeds, enabling research on malware in a safe, controlled environment; *privacy*—ways to share data between organizations in a way that is safe for individual privacy, while providing high research utility; *education*—how to effectively teach security classes and develop hands-on teaching materials for release to a large community of educators via the DeterLab testbed.

ISI recently received an important grant from the Department of Homeland Security to develop inter-network DDoS defenses. The *SENSS Project* is a security service for collaborative mitigation of distributed attacks. An ISP deploying SENSS offers a pay-to-defend service via simple and generic APIs that can be easily implemented in today's routers. Results from this project are being considered for deployment in several large academic and commercial ISPs.

Digital Optical Network Processing and the Optical Turing Machine

High-speed, long-distance communication is necessarily optical, and all-optical digital processing is needed to support forwarding, security, data mining, and filtering on such streams at line rate while in-transit. Today's optical networks, however, treat data communication (transmission of data between sites over a fiber) and packet processing (computations that are performed at specific points within the network to forward and route data) as two logically separate problem domains, using optics for the former and electronics for the latter. This current approach introduces fundamental, unavoidable overheads, significantly limiting network performance and increasing complexity.

ISI researchers have developed the *Optical Turing Machine,* which explores encoding and digital processing mechanisms that unify optical communication and computation by developing optical processing technologies that operate directly on the data formats and encodings used for optical data transmission. Consequently, optical data network processing and data transmission become seamless, enabling dramatically higher-speed, lower-complexity optical networks.

The present work focuses on phase-keyed signals and optical wave mixing to overcome the speed limitations inherent in switching of optical transistor-like devices. Our team of researchers is currently exploring approaches to degenerate phase-sensitive amplification to enable signal regeneration that is compatible with nonlinear device processing for computation.

INFORMATICS RESEARCH SYSTEMS

Innovative Data Management Architectures and Annotation Tools

Researchers from the Informatics Systems Division are working closely with a broad range of domain scientists to create innovative new technology that solves problems of great societal importance. One of our core thrusts is to tackle the issue of increasing volumes and complexity of data and how these pressures can impede the rate of new discovery by scientists, engineers and researchers. To address these issues, we have developed innovative new data management architectures and work closely with scientists to apply these platforms to important problems.

In our collaborative research with the Broad Center for Regenerative Medicine at USC, we are seeking to understand how the structures for the human kidney develop so that we may be better positioned to correct or repair malfunctioning kidneys. In the accompanying figure, we show an interactive annotation tool that was developed using our technology and is now being used by our collaborators to develop an atlas that will track how kidney structures form across different developmental stages.

Interactive annotation tool for tracking kidney structures during developmental stages

Efficient Inner Product Computation for Image and Video Analysis
Shiv Vitaladevuni , Rohit Prasad, and **Prem Natarajan**
US Patent8 953,892, February 2015

Subcircuit Physical Level Power Monitoring Technology for Real-Time Hardware Systems and Simulators
Young Cho and **Siddharth Bhargav**
US Provisional Patent Application 62/170,256, July 2015

Personalized Machine Translation via Online Adaptation
Daniel Marcu and **Jonathan May**
US Patent 9,152,622, July 2015

Translating Documents Based on Content
Daniel Marcu, Radu Soricut and Swamy Viswanathan
US Patent 8,990,064, 2015

Systems & Methods for Identifying Parallel Documents and Sentence Fragments in Multilingual Document Collections
Daniel Marcu and Dragos Munteanu
US Patent 8,943,080, 2015

Segmentation Co-Clustering
Shiv Vitaladevuni, Rohit Prasad, and **Prem Natarajan**
US Patent 9,224,207, December 2015

2015-2016 PUBLICATIONS

COMPUTATIONAL SYSTEMS AND TECHNOLOGY

Reproducibility of Execution Environments in Computational Science Using Semantics and Clouds
Idafen Santana-Perez, **Rafael Ferreira da Silva, Mats Rynge, Ewa Deelman**, María Pérez-Hernández and Oscar Corcho
Journal, Future Generation Computer Systems, 2016

Pegasus in the Cloud: Science Automation through Workflow Technologies
Ewa Deelman, Karan Vahi, Mats Rynge, Gideon Juve, Rajiv Mayani and Rafael Ferreira da Silva
Journal, IEEE Internet Computing, 2016

A Study of Pointer-Chasing Performance on Shared-Memory Processor-FPGA Systems
Gabriel Weisz, Joseph Melber, Yu Wang, Kermin Fleming, Eriko Nurvitadhi and James Hoe
ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA '16), Feb 2016

Panorama: Modeling The Performance of Scientific Workflows
Ewa Deelman, Christopher Carothers, Anirban Mandal, Brian Tierney, Jeffrey Vetter, Ilya Baldin, Claris Castillo, **Gideon Juve**, et al.
Workshop on Modeling and Simulation of Systems and Applications (ModSim), 2015

Performance of the Quantum Adiabatic Algorithm on Constraint Satisfaction and Spin Glass Problems
Itay Hen and A. P. Young
European Physical Journal Special, 2015

Quantum Gates with Controlled Adiabatic Evolutions
Itay Hen
Physical Review, 2015

Fidelity-Optimized Quantum State Estimation
Amir Kalev and **Itay Hen**
New Journal of Physics, 2015

Numerical Studies of the Quantum Adiabatic Algorithm and Spin Glass Problems
Itay Hen and A. P. Young
Proceedings of Conference on Computational Physics (CCP2014), 2015

Independent Functional Testing of Commercial FPGA Devices
Matthew French,Neil Steiner, et al.
GOMACTech, 2015

Gaining Insight into Femtosecond-scale CMOS Effects Using FPGAs
Kenneth Zick, **Matthew French** and S. Li
GOMACTech, 2015

Unraveling Quantum Annealers Using Classical Hardness
Victor Martin-Mayor and **Itay Hen**
Scientific Reports 5, 2015

Probing for Quantum Speedup in Spin Glass Problems with Planted Solutions
Itay Hen, Joshua Job, **Tameem Albash**, Troels Roennow, Matthias Troyer and Daniel Lidar
Physical Review, 2015

Quantum Annealing Correction with Minor Embedding
Walter Vinci, **Tameem Albash**, Gerardo Paz-Silva, **Itay Hen** and Daniel Lidar
Physical Review, 2015

Reexamination of the Evidence for Entanglement in the D-Wave Processor
Tameem Albash, **Itay Hen**, **Federico Spedalieri** and Daniel Lidar
Physical Review, 2015

Quantum Annealing for Constrained Optimization
Itay Hen and **Federico Spedalieri**
Applied Physics Review, 2015

Evaluating High-Level Program Invariants Using Reconfigurable Hardware
Joonseok Park and **Pedro Diniz**
ACM Trans. on Reconfigurable Technology and Systems (TRETS), ACM Press, Nov. 2015

Instrumentation and Mapping Strategies with the LARA Aspect-Oriented Programming Approach
João Cardoso, José Gabriel Coutinho, Tiago Carvalho, **Pedro Diniz**, Zlatko Petrov, Wayne Luk, et al.
Software: Practice and Experience, 2015

An Application-Oriented Cache Memory Configuration for Energy Efficiency in Multi-Cores
Bruno de Abreu Silva, Lucas Cuminato, Alexandre Delbem, **Pedro Diniz** and Vanderlei Bonato
ET Design & Test of Computers, Elsevier, vol. 9, no. 1, 2015

CoRAM++: Supporting Data-Structure-Specific Memory Interfaces for FPGA Computing *(Best Paper Award)*
Gabriel Weisz and James Hoe
25th International Conference on Field Programmable Logic & Applications (FPL), 2015

Non-Intrusive Wireless Links in Commodity Devices
James Couch, Neil Steiner, **Andrew Schmidt**, William Headley, Peter Athanas, Robert McGwier, Alexandra Poetter, S. Rowe and **Matthew French**
Journal of Intelligence Community Research and Development, 2015

Experimental Quantum Annealing: Case Study Involving the Graph Isomorphism Problem
Kenneth Zick, Omar Shehab, **Matthew French**
Scientific Reports, 2015

A Unified Hardware/Software MPSoC System Construction and Run-Time Framework
Sam Skalicky, **Andrew Schmidt**, Sonia Lopez and **Matthew French**
Conference on Design, Automation and Test in Europe (DATE), 2015

A Parallelizing Matlab Compiler Framework and Run Time for Heterogeneous Systems
Sam Skalicky, Sonia Lopez, Marcin Lukowiak and **Andrew Schmidt**
IEEE International Conference on High Performance Computing and Communications (HPCC), 2015

Heterogeneous Cloud Computing: The Way Forward
Stephen Crago and **John Paul Walters**
Computer, vol. 48, no. 1, 2015

Supporting High Performance Molecular Dynamics in Virtualized Clusters Using IOMMU, SR-IOV
Andrew Younge, **John Paul Walters**, **Stephen Crago** and Geoffrey Fox
Proceedings of 11th ACM SIGPLAN/SIGOPS Int'l Conf. on Virtual Execution Environments, 2015

A Game Theoretic Framework for Modeling Adversarial Cyber Security Game Among Attackers, Defenders, & Users
Tatyana Ryutov, **Michael Orosz**, **James Blythe** and Detlof von Winterfeldt
Proceedings of 11th Int'l Workshop on Security and Trust Management (STM), 2015

Maestro Software and Application Performance
John Paul Walters, **David Kang** and **Stephen Crago**
GOMACTech, 2015

Building a Chemical-Protein Interactome on the Open Science Grid
Rob Quick, Soichi Hayashi, Samy Meroueh, **Mats Rynge**, Scott Teige, Bo Wang and David Xu
Proceedings of Science, International Symposium on Grids and Clouds (ISGC), 2015

Storage-aware Algorithms for Scheduling of Workflow Ensembles in Clouds
Piotr Bryk, Maciej Malawski, **Gideon Juve** and **Ewa Deelman**
Journal of Grid Computing, 2015

Adapting Scientific Workflows on Networked Clouds Using Proactive Introspection
Anirban Mandal, Paul Ruth, Ilya Baldin, Yufeng Xin, Claris Castillo, **Gideon Juve**, **Mats Rynge**, **Ewa Deelman**, et al.
IEEE/ACM Utility and Cloud Computing (UCC), 2015

PANORAMA: An Approach to Performance Modeling and Diagnosis of Extreme Scale Workflows
Ewa Deelman, Christopher Carothers, Anirban Mandal, Brian Tierney, Jeffery Vetter, Ilya Baldin, Claris Castillo, **Gideon Juve**, et al.
International Journal of High Performance Computing Applications, 2015

Online Task Resource Consumption Prediction for Scientific Workflows
Rafael Ferreira da Silva, **Gideon Juve**, **Mats Rynge**, **Ewa Deelman** and Miron Livny
Parallel Processing Letters, 25:3, 2015

Practical Resource Monitoring for Robust High Throughput Computing
Gideon Juve, Benjamin Tovar**,** **Rafael Ferreira da Silva**, Daruysz Krol, Douglas Thain, **Ewa Deelman**, et al.
Workshop on Monitoring and Analysis for High Performance Computing Systems Plus Applications, 2015

Characterizing a High Throughput Computing Workload: The Compact Muon Solenoid (CMS) Experiment at LHC
Rafael Ferreira da Silva, **Mats Rynge**, **Gideon Juve**, Igor Sfiligoi, **Ewa Deelman**, et al.
Procedia Computer Science: Int'l Conference On Computational Science (ICCS) —
Computational Science at the Gates of Nature 51, 2015

Reproducibility of Neuroimaging Analyses Across Operating Systems
Tristan Glatard, Lindsay Lewis, **Rafael Ferreira da Silva**, Reza Adalat, Natacha Beck, Claude Lepage, Pierre Rioux, Marc-Etienne Rousseau, Tarek Sherif, **Ewa Deelman**, et al.
Frontiers in Neuroinformatics, 2015

The Case for Resource Sharing in Scientific Workflow Executions
Ricardo Oda, Daniel Cordeiro, **Rafael Ferreira da Silva**, **Ewa Deelman** and Kelly Braghetto
XVI Simposio em Sistemas Computacionais de Alto Desempenho (WSCAD), 2015

Understanding the Scientific Software Ecosystem and its Impact: Current and Future Measures
James Howison, **Ewa Deelman**, Michael McLennan, **Rafael Ferreira da Silva** and James Herbsleb
Research Evaluation, 24:4, 2015

Scheduling Multilevel Deadline-Constrained Scientific Workflows on Clouds Based on Cost Optimization
Maciej Malawski, Kamil Figiela, Marian Bubak, **Ewa Deelman** and Jarek Nabrzyski
Scientific Programming, 2015

Algorithms for Cost- and Deadline-Constrained Provisioning for Scientific Workflow Ensembles in IaaS Clouds
Maciej Malawski, **Gideon Juve**, **Ewa Deelman** and Jarek Nabrzyski
Future Generator Computer Systems 48, 2015

Pegasus: A Workflow Management System for Science Automation
Ewa Deelman, **Karan Vahi**, **Gideon Juve**, **Mats Rynge**, Scott Callaghan, Phil Maechling, **Rajiv Mayani**, Weiwei Chen, **Rafael Ferreira da Silva**, et al.
Future Generation Computer Systems 46, 2015

Using Imbalance Metrics to Optimize Task Clustering in Scientific Workflow Executions
Weiwei Chen, **Rafael Ferreira da Silva**, **Ewa Deelman** and Rizos Sakellariou
Future Generation Computer Systems 46, 2015

Boutiques: An Application-Sharing System Based on Linux Containers
Tristan Glatard, **Rafael Ferreira da Silva**, Nouha Boujelben, Reza Adalat, Natacha Beck, Pierre Rioux, Marc-Etienne Rousseau, **Ewa Deelman,** et al.
NeuroInformatics, 2015

Climate Science Performance, Data and Productivity on Titan
Benjamin Mayer, Patrick Worley, **Rafael Ferreira da Silva** and Abigail Gaddis
Cray User Group Conference, 2015

Panorama: Modeling The Performance of Scientific Workflows
Ewa Deelman, Christopher Carothers, Anirban Mandal, Brian Tierney, Jeffrey Vetter, Ilya Baldin, Claris Castillo, **Gideon Juve**, Dariusz Król, Vickie Lynch, Ben Mayer, Jeremy Meredith, Thomas Proffen, Paul Ruth and **Rafael Ferreira da Silva**
Workshop on Modeling and Simulation of Systems and Applications (ModSim), 2015

Addressing Angular Single-Event Effects in the Estimation of On-Orbit Error Rates
David Lee, Gary Swift, Michael Wirthlin and **Jeffrey Draper**
IEEE Transactions on Nuclear Science, 2015

PathFinder: A Signature-Search Miniapp and its Runtime Characteristics
Aditya Deshpande, **Jeffrey Draper**, J. Brian Rigdon and Richard Barrett
Proceedings of t5th Workshop on Irregular Applications: Architectures and Algorithms, 2015

Congestion-Aware Optimal Techniques for Assigning Inter-Tier Signals to 3D-Vias in a 3DIC
Gopi Neela and **Jeffrey Draper**
Proceedings of 3DIC Conference, 2015

Modeling Data Movement in the Memory Hierarchy in HPC Systems
Aditya Deshpande and **Jeffrey Draper**
Proceedings of MEMSYS Conference, 2015

The Impact of Stress-Induced Defects on MOS Electrostatics and Short-Channel Effects
Ivan Sanchez Esqueda
Journal of Solid-State Electronics, 2015

Compact Modeling of Total Ionizing Dose and Aging Effects in MOS Technologies
Ivan Sanchez Esqueda, Hugh Barnaby and Michael Patrick King
IEEE Transactions on Nuclear Science, vol. 62, no. 4, 2015

The Impact of Defect Scattering on the Quasi-Ballistic Transport of Nanoscale Conductors
Ivan Sanchez Esqueda, Cory Cress, Yu Cao, Yuchi Che and C. Zhou
Journal of Applied Physics, vol. 117, no. 8, 2015

Modeling Radiation Induced Scattering in Graphene
Ivan Sanchez Esqueda and Cory Cress
IEEE Transactions on Nuclear Science, vol. 62, no. 8, 2015

Hydrogen Limits for Total Dose and Dose Rate Response in Linear Bipolar Circuits
Philippe Adell, Bernard Rax, **Ivan Sanchez Esqueda** and Hugh Barnaby
IEEE Transactions on Nuclear Science, vol. 62, no. 6, 2015

INTELLIGENT SYSTEMS

Document Image Quality Assessment Using Discriminative Sparce Representation
Xujun Peng, **Huaigu Cao**, and **Premkumar Natarajan**
12th IAPR Int'l Workshop on Document Analysis Systems (DAS), 2016

Towards a Dataset for Human Computer Communication via Grounded Language Acquisition
Yonatan Bisk, Daniel Marcu and W. Wong
Proceedings of the AAAI 2016 Workshop on Cognitive Symbiotic Systems

Face Recognition Using Deep Multi-Pose Representations
Wael AbdAlmageed, **Yue Wu**, **Stephen Rawls**, Shai Harel, **Tal Hassner**, Iacopo Masi, **Jongmoo Choi**,
Jatuporn Toy Leksut, Jungyeon Kim, **Prem Natarajan**, et al.
Proceedings of IEEE Winter Conference on Applications of Computer Vision (WACV), 2016

The DARPA Twitter Bot Challenge
V.S. Subrahmanian, Amos Azaria, Skylar Durst, Vadim Kagan, **Aram Galstyan**, **Kristina Lerman**, **Linhong Zhu**, **Emilio Ferrara**,
Alessandro Flammini, Filippo Menczer, Rand Waltzman, Andrew Stevens, Alexander Dekhtyar, **Shuyang Gao**, et al.
Journal, IEEE Computer, 2016

BotOrNot: A System to Evaluate Social Bots
Clayton Davis, Onur Varol, **Emilio Ferrara**, Filippo Menczer and Alessandro Flammini
ACM WWW Companion, 2016

2D Sudoku Associated Bijections for Image Scrambling
Yue Wu, Yicong Zhou, Sos Agaian and Joseph P. Noonan
Information Sciences, 327, 2016

Extracting Biomolecular Interactions Using Semantic Parsing of Biomedical Text
Sahil Garg, **Aram Galstyan**, **Ulf Hermjakob** and **Daniel Marcu**
AAAI'16, 2016

Latent Space Models for Multimodal Social Data
Yoon Sik Cho, **Emilio Ferrara**, **Greg Ver Steeg** and **Aram Galstyan**
WWW'16, 2016

Network Structure and Resilience of Mafia Syndicates
Santa Agreste, Salvatore Catanese, Pasquale De Meo, **Emilio Ferrara** and Giacomo Fiumara
Information Sciences 351:30-47, 2016

Dynamically Generated Metadata and Replanning by Interleaving Workflow Generation and Execution
Yolanda Gil and **Varun Ratnakar**
Proceedings of Tenth IEEE International Conference on Semantic Computing (ICSC), 2016

Teaching Big Data Analytics Skills with Intelligent Workflow Systems
Yolanda Gil
Proceedings of the Sixth Symposium on Educational Advances in Artificial Intelligence (EAAI), co-located
with the National Conference of the Association for the Advancement of Artificial Intelligence (AAAI), 2016

The "Majority Illusion" in Social Networks
Kristina Lerman, Xiaoran Yan and **Xin-Zeng Wu**
PLoS ONE, 11(2), 2016

Real-Time Anomaly Detection Framework for Many-Core Routers through Machine Learning Techniques
Amey Kulkarni, Youngok Pino, **Matthew French** and Tinoosh Mohsenin
ACM Journal on Emerging Technologies in Computing (JETC), 2016

Pose-Aware Face Recognition in the Wild
Iacopo Masi, **Stephen Rawls**, **Gérard Medioni** and **Prem Natarajan**
Accepted to IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), 2016

Detection of Promoted Social Media Campaigns
Emilio Ferrara, Onur Varol, Filippo Menczer and Alessandro Flammini
AAAI Int'l Conference on Web and Social media (ICWSM), 2016

Learning the Semantics of Structured Data Sources
Mohsen Taheriyan, **Craig Knoblock**, **Pedro Szekely** and **José Luis Ambite**
Journal, Web Semantics: Science, Services and Agents on the World Wide Web, 2016

Maximizing Correctness with Minimal User Effort to Learn Data Transformations
Bo Wu and **Craig Knoblock**
Proceedings of the International Conference on Intelligent User Interfaces, 2016

Using Syntax-Based Machine Translation to Parse English into Abstract Meaning Representation
Michael Pust, **Ulf Hermjakob**, **Kevin Knight**, **Daniel Marcu** and **Jonathan May**
Proceedings of EMNLP, 2015

How Much Information Does a Human Translator Add to the Original?
Barret Zoph, **Marjan Ghazvininejad** and **Kevin Knight**
Proceedings of EMNLP, 2015

Unifying Bayesian Inference and Vector Space Models for Improved Decipherment
Qing Dou, **Ashish Vaswani**, **Kevin Knight** and Chris Dyer
Proceedings of ACL, 2015

How to Make a Frenemy: Multitape FSTs for Portmanteau Generation
Aliya Deri and **Kevin Knight**
Proceedings of NAACL, 2015

How to Memorize a Random 60-Bit String
Marjan Ghazvininejad and **Kevin Knight**
Proceedings of NAACL, 2015

Unsupervised Entity Linking with Abstract Meaning Representation
Xiaoman Pan, Taylor Cassidy, **Ulf Hermjakob**, Heng Ji and **Kevin Knight**
Proceedings of NAACL, 2015

Building and Using a Knowledge Graph to Combat Human Trafficking *(Best Applied Paper award)*
Pedro Szekely, **Craig Knoblock**, **Jason Slepicka**, Andrew Philpot, Amandeep Dingh, **Chengye Yin**, **Dipsy Kapoor**,
Prem Natarajan, **Daniel Marcu**, **Kevin Knight**, David Stallard, Subessware Karunamoorthy and **RajaGopal Bojanapalli**
14th International Semantic Web Conference (ISWC), 2015

Model Invertibility Regularization (MIR): Sequence Alignment With or Without Parallel Data
Tomer Levinboim, **Ashish Vaswani** and David Chiang
Proceedings of NAACL, 2015

Parsing English into Abstract Meaning Representation Using Syntax-based Machine Translation
Michael Pust, **Ulf Hermjakob**, **Kevin Knight**, **Daniel Marcu** and **Jonathan May**
Proceedings of EMNLP, 2015

A Scalable Architecture for Extracting, Aligning, Linking, and Visualizing Multi-Int Data
Craig Knoblock and **Pedro Szekely**
Proceedings of Conference on the Next Generation Analyst III, 2015

Personalized Trip Planning by Integrating Multimodal User-Generated Content
Om Patri, Ketan Singh, **Pedro Szekely**, Anand Panangadan and Viktor Prasanna
IEEE International Conference on, Semantic Computing (ICSC), 2015

Assigning Semantic Labels to Data Sources
S. K. Ramnandan, Amol Mittal, **Craig Knoblock** and **Pedro Szekely**
Proceedings of the 12th ESWC, 2015

Leveraging Linked Data to Infer Semantic Relations within Structured Sources
Mohsen Taheriyan, **Craig Knoblock**, **Pedro Szekely**, **José Luis Ambite** and **Yinyi Chen**
Proceedings 6th International Work. Consum. Linked Data (COLD), 2015

Mixed Membership Stochastic Blockmodels for the Human Connectome
Daniel Moyer, Boris Gutman, Gautam Prasad, **Greg Ver Steeg** and Paul Thompson
Proceedings of Bayesian and Graphical Imaging for Biomedical Imaging (BAMBI), 2015

Information-Theoretic Clustering of Neuroimaging Metrics Related to Cognitive Decline in the Elderly
Madelaine Daianu, **Greg Ver Steeg**, Adam Mezher, Neda Jahanshad, Talia Nir, Xiaoran Yan, Gautam Prasad,
Kristina Lerman, **Aram Galstyan** and Paul Thompson
Proceedings of the MICCAI Workshop on Medical Computer Vision, 2015

Blockmodels for Connectome Analysis
Daniel Moyer, Boris Gutman, Gautam Prasad, Joshua Faskowitz, **Greg Ver Steeg** and Paul Thompson
11th Int'l Symposium on Medical Information Processing and Analysis, 2015

Estimating Mutual Information by Local Gaussian Approximation
Shuyang Gao, **Greg Ver Steeg** and **Aram Galstyan**
Uncertainty in Artificial Intelligence (UAI), 2015

Understanding Confounding Effects in Linguistic Coordination: An Information-Theoretic Approach
Shuyang Gao, **Greg Ver Steeg** and **Aram Galstyan**
PLoS ONE, 10(6), 2015

Information-Theoretic Characterization of Blood Panel Predictors for Brain Atrophy & Cognitive Decline in the Elderly
Sarah Madsen, **Greg Ver Steeg**, Adam Mezher, Neda Jahanshad, Talia Nir, Xue Hua, Boris Gutman,
Aram Galstyan and Paul Thompson
IEEE International Symposium on Biomedical Imaging, 2015

Maximally Informative Hierarchical Representations of High-Dimensional Data
Greg Ver Steeg and **Aram Galstyan**
Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics (AISTATS), 2015

Efficient Estimation of Mutual Information for Strongly Dependent Variables
Shuyang Gao, **Greg Ver Steeg** and **Aram Galstyan**
Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics (AISTATS), 2015

Disentangling the Lexicons of Disaster Response in Twitter
Nathan Hodas, **Greg Ver Steeg**, Joshua Harrison, Satish Chikkagoudar, Eric Bell and Courtney Corley
3rd International Workshop on Social Web for Disaster Management (SWDM'15), 2015

Black Holes are Almost Optimal Quantum Cloners
Christoph Adami and **Greg Ver Steeg**
Journal of Physics A 58, 2015

Measuring Emotional Contagion in Social Media
Emilio Ferrara and Zeyao Yang
PLoS ONE 10(11), 2015

Trust Networks: Topology, Dynamics and Measurements
Santa Agreste, Pasquale De Meo, **Emilio Ferrara**, Sebastiano Piccolo and Alessandro Provetti
IEEE Internet Computing, 2015

Quantifying the Effect of Sentiment on Information Diffusion in Social Media.
Emilio Ferrara and Zeyao Yang
Peer J. Computer Science 1, 2015

Active Inference for Binary Symmetric Hidden Markov Models
Armen Allahverdyan and **Aram Galstyan**
Journal of Statistical Physics, Volume 161, Issue 2, 2015

Memory-Induced Mechanism for Self-Sustaining Cascades in Networks
Armen Allahverdyan, **Greg Ver Steeg** and **Aram Galstyan**
Journal Physical Review E, 2015

Understanding Confounding Effects in Linguistic Coordination: An Information-Theoretic Approach
Shuyang Gao, **Greg Ver Steeg** and **Aram Galstyan**
PLoS ONE, 10(6), 2015

Human Tutorial Instruction in the Raw
Yolanda Gil
ACM Transactions on Interactive Intelligent Systems, 5(1), 2015

Use of Semantic Workflows to Enhance Transparency and Reproducibility in Clinical Omics
Christina Zheng, **Varun Ratnakar**, **Yolanda Gil** and Shannon McWeeney
Genome Medicine, 7(73), 2015

Field Data Management: Integrating Cyberscience and Geoscience
Matty Mookerjee, Daniel Vieira, Marjorie Chan, **Yolanda Gil,** Terry Pavlis, et al.
Earth and Space Science News, Vol. 96, No. 20, 2015

We Need to Talk: Facilitating Communication Between Field-Based Geoscience & Cyberinfrastructure Communities
Matty Mookerjee, Daniel Vieira, Marjorie Chan, **Yolanda Gil**, Charles Goodwin, Thomas Shipley and Basil Tikoff
GSA Today, Vol. 25, No. 11, 2015

A Task-Centered Interface to On-Line Collaboration in Science
Felix Michel, **Yolanda Gil**, **Varun Ratnakar** and Matheus Hauder
Proceedings of the ACM Conference on Intelligent User Interfaces (IUI), 2015

Supporting Open Collaboration in Science through Explicit and Linked Semantic Description of Processes
Yolanda Gil, Felix Michel, **Varun Ratnakar**, Jordan Read, Matheus Hauder, Christopher Duffy, Paul Hanson and Hilary Dugan
Proceedings of the Twelfth European Semantic Web Conference (ESWC), 2015

A Semantic, Task-Centered Collaborative Framework for Science
Yolanda Gil, Felix Michel, **Varun Ratnakar** and Matheus Hauder
Proceedings of the Twelfth European Semantic Web Conference (ESWC), 2015

A Virtual Crowdsourcing Community for Open Collaboration in Science Processes
Felix Michel, **Yolanda Gil**, **Varun Ratnakar** and Matheus Hauder
Proceedings of the Americas Conference on Information Systems (AMCIS), 2015

A Task-Centered Framework for Computationally Grounded Science Collaborations
Yolanda Gil, Felix Michel, **Varun Ratnakar**, Matheus Hauder, Christopher Duffy and Paul Hanson
Proceedings of the Eleventh IEEE International Conference on eScience, 2015

The Provenance Bee Wiki: Tracking the Growth of Semantic Wiki Communities
Yolanda Gil, **Dipsy Kapoor**, Reed Markham and **Varun Ratnakar**
Proceedings of the Eighth ACM International Conference on Knowledge Capture, 2015

OntoSoft: Capturing Scientific Software Metadata
Yolanda Gil, **Varun Ratnakar** and **Daniel Garijo**
Proceedings of the Eighth ACM International Conference on Knowledge Capture, 2015

Feature Selection using Partial Least Squares Regression and Optimal Experiment Design
Varun Nagaraja and **Wael AbdAlmageed**
International Joint Conference on Neural Networks, 2015

A Graphical Model Approach for Matching Partial Signatures
Xianzhi Du, David Doermann and **Wael AbdAlmageed**
IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), June 2015

Emotion Recognition in the Wild via Convolutional Neural Networks and Mapped Binary Patterns
Gil Levi and **Tal Hassner**
Proc. ACM International Conference on Multimodal Interaction (ICMI), 2015

Effective Face Frontalization in Unconstrained Images
Tal Hassner, Shai Harel, Eran Paz and Roee
IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), 2015

Age and Gender Classification Using Convolutional Neural Networks
Gil Levi and **Tal Hassner**
IEEE Workshop on Analysis and Modeling of Faces and Gestures (AMFG),
at IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), 2015

Schizconnect: Mediating Neuroimaging Databases on Schizophrenia & Related Disorders for Large-Scale Integration
Lei Wang, Kathryn Alpert, Vince Calhoun, Derin Cobia, David Keator, Margaret King, Alexandr Kogan,
Drew Landis, Marcelo Tallis, Matthew Turner, Steven Potkin, Jessica Turner and **José Luis Ambite**
Journal, Neuroimage, 2015

Northwestern U. Schizophrenia Data Sharing for SchizConnect: A Longitudinal Dataset for Large-Scale Integration
Alex Kogan, Kathryn Alpert, **José Luis Ambite**, Daniel Marcus and Lei Wang
Journal, Neuroimage, 2015

Data Integration in the Life Sciences: Lecture Notes in Computer Science
Naveen Ashish and **José Luis Ambite**
11th Int'l Conference (DILS), Springer, 2015

Schizconnect: Virtual Data Integration in Neuroimaging
José Luis Ambite, Marcelo Tallis, Kathryn Alpert, David Keator, Margaret King, et al.
Proceedings of the 11th Int'l Conference on Data Integration in the Life Sciences (DILS), 2015

GEM: The GAAIN Entity Mapper
Naveen Ashish, Peehoo Dewan, **José Luis Ambite** and Arthur Toga
Proceedings of the 11th Int'l Conference on Data Integration in the Life Sciences (DILS), 2015

Terminology Development Towards Harmonizing Multiple Clinical Neuroimaging Research Repositories
Jessica Turner, Danielle Pasquerello, Matthew Turner, David Keator, Kathryn Alpert, Margaret King,
Drew Landis, Vince Calhoun, Steven Potkin, Marcelo Tallis, **José Luis Ambite** and Lei Wang
Proceedings of 11th Int'l Conference on Data Integration in the Life Sciences (DILS), 2015

Leveraging Linked Data to Infer Semantic Relations within Structured Sources
Mohsen Taheriyan, **Craig Knoblock**, **Pedro Szekely**, **José Luis Ambite** and **Yinyi Chen**
Proceedings of the 6th Int'l Workshop on Consuming Linked Data (COLD), 2015

Exploiting Semantics for Big Data Integration
Craig Knoblock and **Pedro Szekely**
AI Magazine, 2015

KR2RML: An Alternative Interpretation of R2RML for Heterogenous Sources
Jason Slepicka, Chengye Yin, **Pedro Szekely** and **Craig Knoblock**
Proceedings of the 6th International Workshop on Consuming Linked Data (COLD), 2015

An Iterative Approach to Synthesize Data Transformation Programs
Bo Wu and **Craig Knoblock**
Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI), 2015

Viral Transcription Alignment
Gil Sadeh, Lior Wolf, **Tal Hassner**, Nachum Dershowitz and Daniel Stokl Ben-Ezr
Int'l Conference on Document Analysis and Recognition (ICDAR), 2015

Teaching Computer Vision: Bringing Research Benchmarks to the Classroom
Tal Hassner and Itzik Bayaz
ACM Transactions on Computing Education (ACM-TOCE), Volume 14 Issue 4, 2015

n–Dimensional Discrete Cat Map Generation Using Laplace Expansions
Yue Wu, Congyun Hua and Yicong Zhou
IEEE Transactions on Cybernetics, 2015

Wide Baseline Stereo Matching with Convex Bounded-Distortion Constraints
Meirav Galun, Tal Amir, **Tal Hassner**, Ronen Basri and Yaron Lipman
Int'l Conference on Computer Vision (ICCV), 2015

Document Image OCR Accuracy Prediction via Latent Dirichlet Allocation
Xujun Peng, **Huaigu Cao** and **Prem Natarajan**
Int'l Conference on Document Analysis and Recognition (ICDAR), 2015

Integrating Natural Language Processing with Image Document Analysis: What We Learned from Two Real-World Applications
Jinying Chen, **Huaigu Cao** and **Premkumar Natarajan**
Int'l Journal of Document Analysis and Recognition (IJDAR), 2015

Boost OCR Accuracy Using iVector Based System Combination Approach
Xujun Peng, **Huaigu Cao** and **Premkumar Natarajan**
SPIE Document Recognition and Retrieval XXII (DRR), 2015

INTERNET AND NETWORKED SYSTEMS

DETERLab and the DETER Project
John Wroclawski, **Terry Benzel**, **Jim Blythe**, Ted Faber, **Alefiya Hussain**, **Jelena Mirkovic** and **Stephen Schwab**
GENI: Prototype of the Next Internet, Springer-Verlag, 2016

Authorization and Access Control: ABAC
Ted Faber, **Steve Schwab** and **John Wroclawski**
GENI: Prototype of the Next Internet, Springer-Verlag, 2016

Mitigating Denial-of-Service Attacks in Wide-Area LQR Control
Nachiappan Chockalingam, Aranya Chakrabortty and **Alefiya Hussain**
Proceedings of forthcoming 2016 IEEE PES General Meeting: Paving The Way for Grid Modernization, 2016

Foundations of Infrastructure CPS
Anuradha Annaswamy, **Alefiya Hussain**, Aranya Chakrabortty and Miloš Cvetković
Proceedings of forthcoming American Control Conference, 2016

DBit: A Methodology for Comparing Content Distribution Networks
Zahib Akhtar, **Alefiya Hussain**, Ethan Katz-Bassett and Ramesh Govindan
Proceedings of Traffic Monitoring and Analysis Workshop (TMA), 2016

All Optical Signal Level Swapping and Multi-Level Amplitude Noise Mitigation using Three Optical Parametric Gain Regions
Yinwen Cao, Morteza Ziyadi, Youichi Akasaka, Amirhossein Mohajerin-Ariaei , Jeng-Yuan Yang, **Joseph Touch**, et. al.
Optics Letters, V41 N4, 2016

Validating Agent-Based Modeling of Human Password Behavior
Bruno Korbar, **Jim Blythe**, Ross Koppel, Vijay Kothari and Sean Smith
Proceedings of AAAI Workshop on Artificial Intelligence for Cyber Security (AICS), 2016

Defending Against Strategic Adversaries in Dynamic Pricing Markets for Smart Grids
Paul Wood, Saurabh Bagchi and **Alefiya Hussain**
Proceedings of 8th Int'l Conference on COMmunication and Systems & NETworkS (COMSNETS), 2016

Accurate In-Situ Runtime Measurement of Energy per Operation of System-on-Chip on FPGA
Si**ddharth Bhargav**, Rishvanth Prabakar and **Young Cho**
Int'l Conference on ReConFigurable Computing and FPGAs (ReConFig), 2015

Accurate Power Measurement Technique for Digital Systems Using Independent Component Analysis
Siddharth Bhargav and **Young Cho**
IEEE Conference on Design of Circuits and Integrated Systems (IEEE DCIS), 2015

T-DNS: Connection-Oriented DNS to Improve Privacy and Security
Liang Zhu, Zi Hu, **John Heidemann**, Duane Wessels, Allison Mankin and Nikita Somaiya
IEEE Security and Privacy, May 2015

When the Internet Sleeps: Correlating Diurnal Networks With External Factors
Lin Quan, **John Heidemann** and **Yuri Pradkin**
ACM Internet Measurements Conference, 2014 (and ongoing data collection)

Different Traffic Models using the LegoTG Framework *(Best Paper Award)*
Genevieve Bartlett and **Jelena Mirkovic**
Int'l Workshop on Computer and Networking Experimental Research Using Testbeds (CNERT), 2015

Cypress: A Testbed for Research in Networked Cyber-Physical Systems
Ryan Goodfellow and **Erik Kline**
Proceedings of 10th EAI Int'l Conference on Testbeds and Research Infrastructures of the Development of Networks & Communities (TRIDENTCOM), 2015

Optimizing Defensive Investments in Energy-Based Cyber-Physical Systems
Paul Wood, Saurabh Bagchi and **Alefiya Hussain**
Proceedings of 20th IEEE Workshop on Dependable Parallel, Distributed and Network-Centric Systems, 2015

Beyond the Lock Icon: Real-time Detection of Phishing Websites Using Public Key Certificates
(Best Paper Honorable Mention)
Zheng Dong, Apu Kapadia, **Jim Blythe** and L. Jean Camp
eCrime Mobile Security Symposium, 2015

Measuring the Security Impacts of Password Policies Using Cognitive Behavioral Agent-Based Modeling
Vijay Kothari, **Jim Blythe,** Ross Koppel and Sean Smith
Proceedings of Symposium and Bootcamp on the Science of Security (HotSoS), 2015

Workarounds to Computer Access in Healthcare Organizations: You want My Password or a Dead Patient?
Ross Koppel, Sean Smith, **Jim Blythe** and Vijay Kothari
Proceedings of Information Technology and Communication in Health Conference (ITCH), 2015

A Roadmap to Evaluate Resilience of Oil and Gas Cyber-Physical Systems
Yatin Wadhawan and **Clifford Neuman**
ACSAC Workshop on Industrial and Control System Security, 2015

Modeling Security Policies for Mitigating the Risk of Load Altering Attacks on Smart Grid Systems
Tatyana Ryutov, Anas AlMajali and **Clifford Neuman**
Proceedings of IEEE Workshop on Modeling and Simulation of Cyber-Physical Energy Systems (MSCPESS), 2015

Mismorphism: a Semiotic Model of Computer Security Circumvention
Sean Smith, Ross Koppel, **James Blythe** and Vijay Kothari
Proceedings of 9th Int'l Symposium on Human Aspects of Information Security & Assurance, 2015

Recommendations for Transport Port Uses
Joseph Touch
RFC 7605 / BCP 165, July 2015, RFC Editor, ISSN 2070-1721

A Widely Deployed Solution to the Generic Routing Encapsulation (GRE) Fragmentation Problem
Ron Bonica, Carlos Pignataro, **Joseph Touch**
RFC 7588, May 2015, RFC Editor, ISSN 2070-1721

Optical Channel De-aggregator of 30-Gbaud QPSK and 20-Gbaud 8 PSK Data using Mapping Onto Constellation Axes
Morteza Ziyadi, Amirhossein Mohajerin-Ariael, Ahmed Almaiman, Yinwen Cao, Mohammad Chitgarha, **Joseph Touch**, et al.
Optics Letters, V40 N21, 2015

Tunable RF Photonics Filter using a Comb-based Optical Tapped-Delay-Line with an Optical Nonlinear Multiplexer
Morteza Ziyadi, Amirhossein Mohajerin-Ariaei, Mohammad Chitgarha, Salman Khaleghi, Ahmed Almaiman, **Joseph Touch**, et al.
Optics Letters, V40 N14, 2015

Phase Noise Mitigation of QPSK Signal Utilizing Phase-Locked Multiplexing of Signal Harmonics and Amplitude Saturation
Amirhossein Mohajerin-Ariaei, Morteza Ziyadi, Mohammad Chitgarha, Ahmed Almaiman, Yinwen Cao, **Joseph Touch**, et al.
Optics Letters, V40 N14, 2015

Experimental Demonstration of Robustness and Accuracy of a DLI-based OSNR Monitor Under Changes in the Transmitter and Link for Different Modulation Formats and Baud Rates
Ahmed Almaiman, Mohammad Chitgarha, Wajih Daab, Morteza Ziyadi, Amirhossein Mohajerin-Ariaei, **Joseph Touch**, et al.
Optics Letters, V40 N9, May 2015, pp. 2012-2015.

Tunable Homodyne Detection of an Incoming QPSK Data Signal Using Two Fixed Pump Lasers
Mohammad Chitgarha, Yinwen Cao, Amirhossein Mohajerin-Ariaei, Morteza Ziyadi, Salman Khaleghi, **Joseph Touch**, et al.
IEEE/OSA Journal of Lightwave Technology, 2015

KESTON EXECUTIVE DIRECTORSHIP ENDOWMENT

USC's Information Sciences Institute has received its first-ever philanthropic gift—the Michael and Linda Keston Executive Directorship Endowment. Under the guidance of the Michael Keston Executive Director Prem Natarajan, the endowment will be used to support the Institute's mission of conducting groundbreaking research in diverse scientific areas such as artificial intelligence, high-performance computing, bioinformatics, information and data management, networking, cyber security, and social analytics. Specifically, the endowment will enable ISI to recruit and fund new, talented, and innovative young researchers–in–residence. These researchers will develop unique new technologies that will see fruition in groundbreaking technologies funded by the U.S. government and industry. They will be aided in their endeavors by ISI's exciting scientific and academic environment, and its dedicated research teams.

Donors Michael and Linda Keston
Michael Keston has a lengthy and distinguished relationship with USC, serving on boards and executive committees and teaching in the university's real estate development program for more than 20 years. A strong supporter of the university, an entrepreneur and a philanthropist, Michael says that he and his wife like to contribute to causes that are important in the world, but that also involve people they themselves would like to work with. His attention was particularly drawn to ISI by the talented and diverse groups of researchers that work here.

Michael Keston Executive Director, Premkumar Natarajan
Dr. Premkumar (Prem) Natarajan, who assumed the leadership of USC's Information Sciences Institute (ISI) in 2013, leads the technical, strategic and operational directions of the institute—a large computer science research facility with campuses in Marina del Rey, CA, and Arlington, VA.

Prior to joining ISI, Prem was with BBN Technologies (later Raytheon BBN Technologies Corp.), an Internet pioneer and subsidiary of the major defense and civilian contractor, Raytheon Company. During his tenure at BBN, he rose through research and leadership ranks to become principal scientist and one of three executive vice presidents.

USC Viterbi
School of Engineering