CERIAS® 2024

Annual Security Symposium ★ April 2nd & 3rd, 2024

Purdue University, West Lafayette, Indiana

Table of Contents

#CERIAS

Opening Keynote

9:00a

STEW 214

"Emerging and Future Cybersecurity Risks for National Security"

Dr. Daniel "Rags" Ragsdale
Deputy Assistant National Cyber Director, White House Office of the National Cyber Director

Dr. Daniel "Rags" Ragsdale is the former Principal Director for Cyber, in the Office of the Director of Defense Research and Engineering (Research & Technology). In that role, Dr. Ragsdale is responsible for coordinating cyber modernization efforts across the Department of Defense, with specific responsibility for the establishment of policies and for supervision over cyber modernization research and engineering, technology development, prototyping, experimentation, developmental testing, and transition activities. He is also responsible for making recommendations concerning the allocation of resources and alignment of efforts across the Department.

Before his recent return to service in the Department of Defense, Dr. Ragsdale was the founding director of the Texas A&M Cybersecurity Center and a Professor of Practice in the Department of Computer Science and Engineering. In the Director role, Dr. Ragsdale was responsible for leading, coordinating, and facilitating cybersecurity research and educational activities across the University. Dr. Ragsdale had previously served as a Program Manager in the Defense Advanced Research Projects Agency (DARPA). In that capacity, he successfully led and managed a $175M research and development portfolio of classified and unclassified cybersecurity and educational programs.

Prior to joining DARPA, Colonel (retired) Ragsdale served 30 years in the U.S. Army in a wide array of operational, educational, and research and development roles. He participated in Operations Urgent Fury (Grenada), Enduring Freedom (Afghanistan) and Iraqi Freedom (Iraq). Dr. Ragsdale, also served nearly 15 years at the United States Military Academy, West Point, in a variety of teaching and research roles. He culminated his Army service as the Vice Dean for Education, the Principal Deputy to West Point's Chief Academic Officer.

Dr. Ragsdale is a 1981 graduate from the US Military Academy. He has earned a Master of Science degree, in Computer Science, from the Naval Postgraduate School and a Ph.D., in Computer Science, from Texas A&M University.

Dr. Ragsdale is a recipient of the Colloquium for Information System Security Education (CISSE) Founder's Medal, the International Federation for Information Processing (IFIP) Outstanding Service Award, the Federal Information Systems Security Education Association (FISSEA) Educator of the Year Award, the Texas A&M Engineering Genesis Award for Multidisciplinary Research, and the US Military Academy Apgar Award for Excellence in Teaching. Among his military awards and decorations are the Secretary of Defense Medal for Outstanding Public Service, the Legion of Merit, six Meritorious Service Medals and the Bronze Star.

Dr. Ragsdale and Cynthia, his wife of thirty-seven years, raised three children. They are each currently serving in the armed force as junior Officers; one in the Army, one in the Air Force, and one in the Navy

Networking Break

10:00a

Lightning Talk

10:15a

"Strategies for Launching Successful Data Poisoning Attacks on Machine Learning Models"

Dr. Jing Gao,
Associate Professor, Elmore Family School of Electrical and Computer Engineering
Purdue University

Abstract

The past decades have witnessed the astonishing predictive power of various machine learning models for different applications, but the usual assumption is that these models are deployed in a benign environment. Despite much progress, the robustness of many machine learning models under attacks remains unexplored. In this talk, I will introduce our work on data poisoning attacks, which aim to manipulate the behavior of machine learning models by injecting malicious or misleading data (i.e., adversarial samples) into the training set. I will present the methodologies we developed to systematically create adversarial samples that could skew the model's original objectives for a series of machine learning tasks, including outcome interpretation, fair machine learning, next-item recommendation, and knowledge graph embedding. Experiments and theoretical analysis of our designed attacks provide valuable insights into the vulnerabilities of machine learning models and could inform the development of more secure and reliable machine learning systems.

About Dr. Gao

Dr. Jing Gao is an Associate Professor in the Elmore Family School of Electrical and Computer Engineering of Purdue University. Before joining Purdue in January 2021, She was an Associate Professor in the Department of Computer Science and Engineering at the University at Buffalo, State University of New York. She received her PhD from Computer Science Department at University of Illinois at Urbana Champaign in 2011 under the supervision of Prof. Jiawei Han. She received M.E. and B.E. from the Computer Science and Technology Department at Harbin Institute of Technology in China.

Panel Discussion #1 (Day 1)

10:35a

"The Security and Trustworthiness of AI in the era of LLM"

STEW 214

Moderator - Dr. Greg Shannon, Laboratory Fellow, Idaho National Laboratory

Dr. Greg Shannon is the Chief Cybersecurity Scientist for Idaho National Laboratory's National and Homeland Security Directorate and is the Chief Science Officer for the Cybersecurity Manufacturing Innovation Institute (CyManII) based at the University of Texas at San Antonio. His research focuses on applications of formal methods to ensure security and resilience properties in cyber-physical systems such as energy-intensive manufacturing and critical infrastructure. To promote more structured awareness of vulnerabilities in cyber-physical systems, he is a co-chair of the newly established CWE-CAPEC special interest group for Industrial Control Systems and Operational Technology. Greg is a member of the U.S. Air Force Science Advisory Board, is a founding board member for Women in Cybersecurity, and has served as the Assistant Director for Cybersecurity Strategy at the White House Office of Science and Technology Policy. Previous to joining INL, he was the Chief Scientist for the CERT Division in the Software Engineering Institute at Carnegie Mellon University. He received a BS in Computer Science from Iowa State University and earned a PhD in Computer Sciences at Purdue University.

Dr. Evercita Eugenio, Statistician, Cyber Operations Research Engineering Group, Sandia National Laboratories

Evercita Eugenio is a statistician with the Cyber Operations Research Engineering group at Sandia National Laboratories in Livermore, CA. Evercita's research is focused on differential privacy, privacy enhancing technologies and machine learning. She is also interested in biostatistics, clinical trials and likelihood theory. Evercita received a PhD in Applied and Computational Mathematics and Statistics from the University of Notre Dame where her dissertation focused on differentially private data synthesis methods. In addition, Evercita has a MS from Oregon State University and BS from the University of Washington.

Dr. Nate Gleason, Program Leader, Cyber and Infrastructure Resilience, Global Security, Lawrence Livermore National Laboratory

Dr. Nate Gleason is the Program Leader for Cyber and Infrastructure Resilience (CIR) within the Energy and Homeland Security Program in the Global Security Directorate at Lawrence Livermore National Laboratory (LLNL).

Gleason joined LLNL as the founding Program Leader for CIR in January 2016 and has since helped the program grow. As Program Leader for Cyber and Infrastructure Resilience, Gleason is responsible for developing technologies and solutions that will allow the nation to progress towards a future where our critical infrastructure systems are intelligent, self-healing, and resilient to cyber and physical disruptions. CIR's mission also works to derive 100 percent of our nation's energy from renewable sources, eliminating its negative impact on the environment as well as reliance on foreign entities for fuel.

Prior to joining LLNL in January 2016, Gleason spent 12 years at Sandia National Laboratories in a variety of technical and management positions including Deputy to the Vice President, Deputy Program Director for Sandia's Homeland Security Program, Department Manager for the Advanced Systems Engineering and Deployment Department and the Systems Research and Analysis Department.

Dr. Julia Rayz, Professor and Associate Department Head, Computer and Information Technology, Purdue University

Dr. Rayz' primary research interests lie in natural language understanding, knowledge discovery and representation, and computational recognition of salient information in texts, as well as in uncertainty management. These belong in the domain of Artificial Intelligence, in the areas of Natural Language Processing, and Cognitive Science on the one hand, and Imprecision Management on the other.

Her long-term research interest and goal are to enable people to communicate with computers informally, using (eventually, any) natural language, with the full understanding of what is said, and perhaps what is more important, of what is left unsaid. While knowledge representation, reasoning, machine learning and computational linguistics are not new areas and have received considerable attention, combining the meaning extracted from natural language texts with our knowledge of the world, represented in some conceptual form, with built-in fuzziness, vagueness, and uncertainty, where necessary, while remaining factually correct when desired—this computational task has not been fully resolved. Her current research attempts to come closer to an understanding of how to construct such a model, and modeling and detecting humor has turned out to be a convenient and visible entry into it as well as a good testing mechanism.

Dr. Lin Tan, Mary J. Elmore New Frontiers Professor of Data Science & Professor of Computer Science, Purdue University

Lin Tan's research interests include software dependability, software-AI synergy, and software text analytics. Some of her research focuses are leveraging machine learning and natural language processing techniques to improve software dependability, and using software approaches to improve the dependability of machine learning systems. Prior to joining Purdue, she was a Canada Research Chair and an associate professor at the University of Waterloo.

PARI

STEW 214

"Introducing Purdue Applied Research Institute (PARI)"

11:35am

Mark Lewis, President and CEO
Purdue Applied Research Institute

Dr. Mark J. Lewis is president and chief executive officer of the Purdue Applied Research Institute (PARI), the nonprofit applied research arm of Purdue University with a particular focus on national security, economic security and food security for the United States. A renowned researcher, professor and former deputy undersecretary of defense, Lewis brings a wealth of national security, scientific and academic experience to the institute.

Lewis came to Purdue from his post as executive director of the National Defense Industrial Association's Emerging Technologies Institute, a nonpartisan think tank focused on technologies that are critical to the future of national defense. This institute provides research and analyses to inform the development and integration of emerging technologies into the defense industrial base.

Before this, Lewis was director of defense for research and engineering in the Defense Department, overseeing technology modernization for the services and defense agencies, as well as the acting deputy undersecretary of defense for research and engineering. In that role, he was the Pentagon's senior-most scientist, managing a $17 billion budget that included the Defense Advanced Research Projects Agency, the Missile Defense Agency, the Defense Innovation Unit, the Space Development Agency, Federally Funded Research and Development Centers (FFRDC) and the Defense Department's basic and applied research portfolio.

From 2012 to 2019, Lewis was the director of the Science and Technology Policy Institute, an FFRDC that supported the Executive Office of the President and other executive branch agencies in forming national science and technology policy.

Lewis is a professor emeritus at the University of Maryland, where he served as the Willis Young Jr. Professor and chair of the Department of Aerospace Engineering until 2012. A faculty member at Maryland for 25 years, he taught and conducted basic and applied research in hypersonic aerodynamics, advanced propulsion and space vehicle design and optimization. Best known for his work in hypersonics, Lewis' research has spanned the aerospace flight spectrum from the analysis of conventional jet engines to entry into planetary atmospheres. From 2004 to 2008, Lewis was the Air Force's chief scientist, the principal scientific adviser to the chief of staff and secretary of the Air Force. As the longest-serving chief scientist in Air Force history, his primary areas of focus included hypersonics, space launch, energy, sustainment, advanced propulsion, basic research and workforce development. From 2010 to 2011, he was president of the American Institute of Aeronautics and Astronautics.

Lewis attended the Massachusetts Institute of Technology, where he received his bachelor of science in aeronautics and astronautics, bachelor of science in Earth and planetary science (1984), and master of science (1985) and doctor of science (1988) in aeronautics and astronautics. He is the author of more than 320 publications and has advised more than 60 graduate students. He has served on boards for NASA and the Defense Department, including two terms on the Air Force Scientific Advisory Board.

A recipient of the Air Force Exemplary, Meritorious and Exceptional Civilian Service Awards and of the Secretary of Defense Outstanding Public Service Award, Lewis was also the 1994 AIAA National Capital Young Scientist/Engineer of the Year; received the IECEC/AIAA Lifetime Achievement Award, the AIAA Dryden Lectureship Award, and the AFA Theodore von Karman Award; and is an Aviation Week and Space Technology Laureate. He is a member of the International Academy of Astronautics, a fellow of the American Society of Mechanical Engineers, a fellow of the Royal Aeronautical Society and an honorary fellow of the American Institute of Aeronautics and Astronautics.

Lightning Talk

12:05p

STEW 214

"Security of AI-enabled CPS and Leveraging LLM to Improve AI-enabled CPS Security"

Dr. Ashok Vardhan Raja, Assistant Professor of Cybersecurity, Computer Information Technology and Graphics (Purdue Northwest)

Abstract

This talk focuses on the safe integration of Artificial Intelligence (AI) with Cyber Physical Systems (CPS), specifically about unmanned aerial vehicles (UAVs). Because they combine physical components and computer algorithms, UAVs are considered a CPS. UAVs have seen significant expansion in a variety of missions in recent years, including infrastructure inspection, thanks to their great mobility and sophisticated sensing capabilities. The potential to integrate AI with CPS is made possible by the development of AI algorithms and technology. Due to potential vulnerabilities in the underlying AI models, this integration raises new security and safety issues. Malicious actors may exploit these weaknesses causing major security and safety concerns. Therefore, it is imperative to guarantee the safe integration of AI and UAVs while bolstering their resistance against hostile settings. To find potential vulnerabilities and related countermeasures, we will first examine the data sensing and processing pipeline of important sensors used in AI-enabled UAV operations in this presentation. The talk will also cover the use of Large Language Models to improve this integration's security.

About Dr. Raja

Ashok Vardhan Raja is an Assistant Professor of Cybersecurity in the department of Computer Information Technology and Graphics for the College of Technology at Purdue University Northwest. His research is on secure integration of Artificial Intelligence (AI) and Cyber Physical Systems (CPS) such as UAVs for robust operations. He is expanding his current work by using Swarm of UAVs to address security issues and to other domains in the integration of AI and CPS.

Lunch Break

12:25p - 1:30p

PMU

Purdue Memorial Union's Atlas Family Marketplace has a wide variety of dining options to satisfy you.

Aatish - Halal contemporary kitchen

BBQ District - Slow-cooked meats, regional sauces and savory sides.

Chef Bill Kim's - Asian dumplings and bowls using authentic ingredients.

Fresh Fare - Fresh flavors with an emphasis on dairy-free and gluten-free options.

Latin Inspired - South American flavors and Latin fare from Brazil and Argentina.

Pizza & Parm Shop - Detroit-style pizza with a caramelized cheese crust and creative parm sandwiches.

Sol Toro - Mexican flavors with a modern flair.

Starbucks®

Sushi Boss - Fresh custom sushi.

Walk On's Sports Bistreaux - Louisiana-inspired cuisine with a game-day flair: burgers, wraps, salads, seafood specialties.

Zen - Build-your-own sushi in a bowl, salad bar and boba teas.

Also... see page 37 for a map of nearby restaurants

TechTalk

1:30p
STEW 214

"How Cyber is Revolutionizing Contemporary Conflict"

Christopher Cleary, Vice President, Global Cyber Practice Innovation & Capabilities Organization, ManTech

Abstract

Mr. Cleary will transport attendees through the annals of military history to understand how the evolution of warfare parallels the rise of new technologies. Drawing from the transformative impact of innovations like the submarine and bomber during World War II, learn not only how these advancement reshaped the modern battlefield, but how senior leadership struggled to fully understand or appreciate how these new capabilities would impact the the outcome of the war. Come listen to how cyber capabilities are similarly revolutionizing contemporary conflict and what part you can play to shape this new domain. Secure, Survive, Strike.

About Chris

Christopher "Chris" Cleary is the Vice President of ManTech's Global Cyber Practice. He directs the technical strategy to expand and enhance ManTech's cognitive cyber innovation and capabilities business.

He joined ManTech after previously serving as The Department of the Navy's Principal Cyber Advisor and Chief Information Security Officer, respectively.

Mr. Cleary has deep expertise in the public and private sectors and the military with more than 25 years of experience driving enterprise Cyber and Signals Intelligence (SIGINT) innovation. He has a proven track record of accelerating technology innovation, driving revenue growth and increasing market expansion for multi-billion-dollar corporations including Leidos, Sparta (now Parsons) and Verizon.

Mr. Cleary has hands-on experience managing critical Cyber programs for U.S. Cyber Command, the Intelligence Community and other key branches of the government.

He is a graduate of the U.S. Naval Academy with a Bachelor of Science in history and a Master of Arts in National Security and Strategic Studies from the U.S. Naval War college.

Panel Discussion #2

2:15p

STEW 214

"Cyber-physical - Secure Lessons Learned for Different Domains"

Moderator - Dr. Dongyan Xu, Director, CERIAS, Purdue University

Randall Brooks
Principal Technical Fellow, RTX

Randall Brooks is a Principal Technical Fellow for RTX (NYSE: RTX). He is the Chief Engineer of the RTX Cyber Operations, Development and Evaluation (CODE) Center, which focuses on product cybersecurity. Randall represents the company within the US International Committee for Information Technology Standards Cyber Security 1 (CS1) and the Cloud Security Alliance (CSA). He has more than 25 years of experience in cybersecurity with a recognized expertise in software assurance (SwA) and secure development life cycles (SDLCs). In addition to holding eight patents, Randall is a CISSP, CSSLP, ISSEP, ISSAP, ISSMP, and CCSK. He graduated from Purdue University with a bachelor's degree from the School of Computer Science.

Doug Kiehl
Senior Director at Eli Lilly and Company

Doug Kiehl is a Senior Director at Eli Lilly and Company and leads the Disruptive/Transformative Technologies Team (DT3) and Digital Twin Center of Excellence with focus on digital transformation, automation, extractables/leachables and next-gen bioprocess. He serves as a member of the United States Pharmacopeia (USP) Packaging and Distribution Expert Committee, Chair for the Product Quality Research Institute (PQRI) Steering Committee, PhRMA Topic Lead for the International Council for Harmonisation (ICH) Q3E Guidance Expert Working Group, Board of Directors for the Extractables/Leachables Safety Information Exchange (ELSIE) Consortium, Chair for the International Society for Optics and Photonics (SPIE) Defense and Commercial Sensing Conference, Executive Governing Council for the Nano-Bio Materials Consortium (NBMC/SEMI) and founding member of the Biomolecule Reactivity Consortium. He has published his work in several peer-reviewed and trade journals, and has organized, chaired and presented at numerous conferences.

Dave Rottier, Chief Engineer, System Software Technology and Embedded Product Cybersecurity, Caterpillar, Inc.

Dave Rottier is the Chief Engineer for System Software Technology and Embedded Product Cybersecurity in Integrated Components & Solutions at Caterpillar, Inc. He leads a global team of engineering fellows, engineering managers, and engineers responsible for developing strategies and processes for embedded product cybersecurity, platform software, common service and network solutions, system diagnostics, and development and process tools.
Dave has over 25 years of experience at Caterpillar and started his career as a leader within electronics for engine systems and machine systems. His experience in electronic systems led to advanced opportunities in information systems including telematics, displays, and embedded cybersecurity. Dave was named Chief Engineer of Embedded Product Cybersecurity in 2021.

Dave earned a Bachelor of Science in electrical engineering from Michigan Technological University, holds two patents, and is a certified DMAIC and DMEDI Black Belt.

Dr. Michael Sangid, Professor of Aeronautics and Astronautics/Professor of Materials Engineering, Purdue University

Michael D. Sangid received his B.S. (2002), M.S. (2005), PhD (2010) in Mechanical Engineering from the University of Illinois at Urbana-Champaign (UIUC). After his Master's degree, Dr. Sangid spent two years working in Indianapolis, IN for Rolls-Royce Corporation. Dr. Sangid is a professor at Purdue University in the School of Aeronautics and Astronautics with a courtesy appointment in Materials Engineering, where he continues his work on building computational materials models for failure of structural materials with experimental validation efforts focused at characterization of the stress/strain evolution at the microstructural scale during in situ loading. He is a recipient of the ASME Orr, TMS Early Career Faculty Fellow, NSF CAREER, and the AFOSR, ONR, and DARPA Young Investigator/Faculty Awards. He is currently serving as an editor of the International Journal of Fatigue. Dr. Sangid has started and serves as the Executive Director of the Hypersonics Advanced Manufacturing Technology Center, which is the first contract of the Purdue Applied Research Institute. Dr. Sangid is also the director of the Purdue Institute of National Security.

Networking Break

3:15p
STEW 214

Fireside Chat

3:30p
STEW 214

Dr. Eugene Spafford, Executive Director Emeritus & Founder, CERIAS, Purdue University

Eugene H. Spafford is a professor of Computer Sciences at Purdue University, a professor of Philosophy (courtesy appointment), and is Executive Director Emeritus of the Center for Education Research Information Assurance and Security. CERIAS is a campus-wide multi-disciplinary Center, with a broadly-focused mission to explore issues related to protecting information and information resources. Spaf has written extensively about information security, software engineering, and professional ethics. He has published over 100 articles and reports on his research, has written or contributed to over a dozen books, and he serves on the editorial boards of most major infosec-related journals.

Dr. Heidi Shyu, Under Secretary of Defense for Research and Engineering, U.S. Dept. of Defense

Ms. Heidi Shyu is the Under Secretary of Defense for Research and Engineering (OUSD(R&E)). In this role, she serves as the Chief Technology Officer for the Department of Defense (DoD), mandated with ensuring the technological superiority of the U.S. military, and is responsible for the research, development, and prototyping activities across the DoD enterprise. She also oversees the activities of the Defense Advanced Research Projects Agency (DARPA), the Missile Defense Agency (MDA), the DoD Laboratory and Engineering Center enterprise, and the Under Secretariat staff focused on developing advanced technology and capability for the U.S. military.

Previously, she served as the Assistant Secretary of the Army for Acquisition, Logistics and Technology (ASA (ALT)), from September 2012 to January 2016. Prior to this, she was Acting ASA (ALT) beginning in June 2011 and appointed the Principal Deputy in November 2010. As the ASA (ALT), she served as the Army Acquisition Executive, the Senior Procurement Executive, the Science Advisor to the Secretary of the Army, and the Army's Senior Research and Development official. She had principal responsibility for all Department of the Army matters related to logistics. Ms. Shyu also led the execution of the Army's acquisition function and the acquisition management system. Her responsibilities included providing oversight for the life cycle management and sustainment of Army weapons systems and equipment from research and development through test and evaluation, acquisition, logistics, fielding, and disposition.

Prior to her government service, Ms. Shyu was the Vice President of Technology Strategy for Raytheon Company's Space and Airborne Systems.

Ms. Shyu holds a Bachelor of Science in mathematics from the University of Brunswick in Canada, a Master of Science degree in mathematics from the University of Toronto, and a Master of Science degree in Electrical Engineering with a focus on System Sciences along with the Engineer's Degree from UCLA. She received an Honorary Doctorate of Science from the University of New Brunswick. She is also a graduate of the UCLA Executive Management Course Program.

A member of the Air Force Scientific Advisory Board from 2000 to 2010, she served as the Vice Chair from 2003 to 2005 and Chair from 2005 to 2008. Ms. Shyu is a member of the National Academy of Engineering and an American Institute of Aeronautics and Astronautics (AIAA) Honorary Fellow.

Lightning Talk

4:15p

STEW 214

"Debugging and Explaining Unfairness in Machine Learning Models"

Dr. Romila Pradhan, Assistant Professor, Computer & Information Technology, Purdue University

Abstract
Algorithmic decision-making systems are increasingly being used to automate consequential decisions in many high-stakes application domains. There is a growing concern that these systems are not transparent, and perpetuate systemic biases reflected in training data. Such discriminatory outcomes violate human rights and undermine public trust in automated decision-making systems. To render these systems more explainable and trustworthy, I will describe our efforts toward making the decisions of AI-based systems less biased. I will present Gopher, a causal data-based explanation mechanism, which allows us to diagnose the outcomes of a machine learning model and detect subsets of the training data most responsible for biased decisions. Gopher quantifies and efficiently approximates the causal responsibility of training data subsets and prunes the huge search space to generate compact, interpretable, and causal explanations for biased model predictions. Experimental evaluation over several real-world datasets in the fair ML literature demonstrates that our system is effective and efficient in generating interpretable explanations for debugging root causes of model bias.

About Dr. Pradhan
Dr. Pradhan is an Assistant Professor in the Department of Computer & Information Technology at Purdue University in West Lafayette, Indiana. Her research interests are in the areas of data management and data science. Her research is driven by the need to build trustworthy and responsible decision-making systems. More recently, Dr. Pradhan has been building systems that facilitate explainability, fairness, and accountability of data-driven decision-making systems.

Before joining Purdue CIT, Dr. Pradhan held a Postdoctoral Researcher position in the Halıcıoğlu Data Science Institute at the University of California San Diego, and was a Visiting Assistant Professor in the Department of Computer Science at Purdue University. She has a Ph.D. in Computer Science from Purdue University, and a B.S. and M.S. in Mathematics and Computing from the Indian Institute of Technology (IIT) Kharagpur, India.

Course Preview

4:35p

STEW 214

"Cybersecurity Awareness and Strategies to Enhance Resilience of Recovery and Response Operations During Disaster"

Dr. Umit Karabiyik, Associate Professor, Computer and Information Technology, Purdue University

MGT 480-W/ MGT 492-W
The Internet of Things (IoT) has become an integral part of our country? s core infrastructure at all levels. This raises the question: Are your response and recovery operations ready for cyber attacks? This course utilizes the whole-of-community approach to increase the awareness of opportunistic IoT based cyber attacks that might occur concurrently during an ongoing disaster response / recovery operation.

In this course, management and incident management teams will learn about the essential elements of cybersecurity and best practices, when to engage information technology, and how to implement strategies during physical disasters, including awareness of the vulnerabilities of IoT-related cyber attacks. Participants will also learn about cybersecurity resilience strategies to mitigate risk and disruption, while maintaining the integrity of the ongoing non-cyber related incident response and recovery operations.

The course uses state of the art artificial intelligence (AI) to adapt the content based on participants? engagement. Course progress is tracked, allowing participants to stop and then continue where they left off. Once enrolled, participants will have 120 days to complete the training at their own pace. Note that participants will be able to instantly switch back and forth from/to the 3D VR based course environment to the web based course environment with the same content during this course

For more information email: CRAVRE@purdue.edu

Research Poster Session Preview

4:45p
STEW 214

Students present their research in an "elevator pitch" format

End of Session 1

5:25p

Poster Session

Convergence Center Lobby 6:30-8:30pm

Highlighting research conducted by students

Day 2

Registration / Coffee

8:00a
STEW 214

Opening Comments, CERIAS Awards

8:45a
STEW 214

Research Poster Awards
We announce winners from this year's poster session.

Diamond Award
The annual Diamond Award goes to a student that most exemplifies the "diamond in the rough" transition through outstanding academic achievement and/or research excellence.

Keynote

9:00a
STEW 214

"Cyberspace, Cybersecurity, and a New World Order"

Samuel Visner, Technical Fellow
Aerospace Corporation

Abstract
While our focus on cybersecurity tends to be technical, to concentrate on adversary tactics, techniques, and procedures, and development of effective cyber defenses, other vital questions loom. Among them:

How did cyberspace evolve? How did that evolution shape cybersecurity? How do members of the international system wield cybersecurity in pursuit of their interest? How is the international system being reshaped, and what do these changes mean for the role of the United States in an increasingly competitive world?

Samuel Sanders Visner, Chair of the Board of Directors of the Space Information Sharing and Analysis Center and adjunct professor at Georgetown University' Program in Science and Technology in International Affairs, will offer a framework that unifies our understanding of cyberspace, cybersecurity, and the international system Using insights gained in the public, private, and academic sectors, Sam will discuss how the geopolitical environment and seek to engage conference participants in a discussion about how US interests can be managed as a changed international system develops.

About Samuel
Samuel Sanders Visner serves as Chair, Board of Directors, Space Information Sharing and Analysis Center and is a Technical Fellow at the Aerospace Corporation. Sam has served as Director of the National Cybersecurity Federally Funded Research and Development Center, General Manager of two cybersecurity businesses, and Chief of Signals Intelligence Programs at the National Security Agency. Sam established at Georgetown University the cybersecurity policy, operations, and technology curriculum, and serves on the Board of Directors of Oak Ridge Associated Universities and as a consultant to the Army Science Board. Sam has a BS degree in International Politics from Georgetown University and an MA in Telecommunications from the George Washington University. The views expressed in Sam's presentation represent are his own.

Networking Break

10:00a
STEW 214

Tech Talk

10:15a

STEW 214

"The Mythical LLM-Month"

Dr. Claudionor N. Coelho Jr, Chief AI Officer, Zscaler, Inc.

Abstract

No doubt. Last year was the year of the LLM. As we move more and more into 2024, we are seeing that almost no LLM promised last year was delivered, or when delivered, they quickly made the headlines.

We are going to show that what people have been claiming to be standalone LLMs are in fact complex software systems where LLMs are an integral part of it. We call these systems AI Agents. A lot of solutions built are still in their infancy, and the race to get these systems up and running caused a huge impact in readiness of such systems. For example, one important aspect of AI Agents to increase their reliability is grounding, which is often forgotten in examples and demos. Finally, most people still regard data as a second class citizen, and in reality for AI Agents based on Retrieval Augmented Generation (or RAG for short), having good data is as important or more important than a LLM.

We will finalize this presentation showing how to merge software engineering practices from the last 50 years into this new era, where multimodal data and algorithms are integrated into vector databases, and used to create high quality copilot applications.

About Dr. Coelho

With nearly three decades in the IT and software engineering industry, Claudionor has extensive research and development experience in machine learning and deep learning techniques (including Generative AI), software systems, cybersecurity and semiconductors. As the Chief AI Officer at Zscaler, he is chartered with driving the vision and implementation of advanced AI technologies to strengthen the world's largest security platform and propel Zscaler's innovation engine forward.

Prior to joining Zscaler, Claudionor served as the Chief AI Officer and SVP of Engineering at Advantest, where he spearheaded the development of a Zero Trust private cloud solution tailored for the semiconductor manufacturing market, enabling it to run Machine Learning workloads. Before that, Claudionor was the VP/Fellow of AI and the Head of AI Labs at Palo Alto Networks where he led the charge in AI, AIOps and Neuro-symbolic AI, an advanced form of AI that enables reasoning, learning, and cognitive modeling, to help revolutionize time series analysis tools on a massive scale. Claudionor's career also includes vital roles in Machine Learning and Deep Learning at Google, where he developed a state-of-the-art Deep Learning technology designed for automatic quantization and model compression which played a pivotal function in the search for subatomic particles at CERN. This work was featured in the cover page of Nature Machine Intelligence in August, 2021.

Claudionor holds a PhD in Electrical Engineering and Computer Science from Stanford University and an MBA from Ibmec in Brazil. He earned both his M.Sc in Computer Science and B.S. in Electrical Engineering from the Universidade Federal de Minas Gerais and is an Invited Professor in the Electrical and Computer Engineering Department at Santa Clara University.

Lightning Talk

11:00a
STEW 214

"Three-body Problem in Privacy Protection: Chaos or New Hope?"

Dr. Wenhai Sun, Assistant Professor, Computer and Information Technology, Purdue University

Abstract
Local differential privacy (LDP) has been widely integrated into commercial use for privatized data collection and analytical tasks. The beauty of LDP not only comes from its rigorous privacy guarantee but also its elegance in an adjustable balance between privacy and utility to satisfy various user and application demands. However, when putting LDP in a real-world scenario, the privacy-utility trade-off becomes fragile when an attacker aims to alter data utility by attempting to manipulate the LDP result. With this new security perspective, we have a new three-body problem in LDP, i.e., the relationships among security, privacy, and utility are uncertain. The consequences are profound from discouraging adoption of the privacy-friendly technologies to harming Internet freedom by suppressing the voice of target groups. In this talk, I will briefly introduce our exploratory work on the contribution to the understanding of the security dimension of LDP and my ongoing NSF CAREER project on a new hope of leveraging machine intelligence to handle the complexity and eventually creating accountable, transparent, and user-friendly "AI for privacy".

About Dr. Sun
Dr. Wenhai Sun is an Assistant Professor in the Department of Computer and Information Technology at Purdue University. He holds two PhDs from the Department of Computer Science at Virginia Tech and the School of Telecommunications Engineering at Xidian University. His research interest mainly lies in designing and developing next-generation secure and accountable networked systems that prioritize the privacy and utility needs of users by leveraging interdisciplinary research including AI/ML, cryptography, trusted hardware, and software engineering. He has published papers in top security, AI, and networking conferences and prestigious journals, such as USENIX Security, NeurIPS, IEEE INFOCOM, IEEE TIFS, IEEE TDSC, IEEE TPDS, and IEEE TSC. He is the recipient of the prestigious NSF CAREER Award. He won the Distinguished Paper Award in ACM ASIACCS 2013. Dr. Sun is a Senior Member of IEEE.

CERIAS Tech Talk

11:20a
STEW 214

"Valid Statistical Inference for Privatized Data"

Dr. Jordan Awan, Assistant Professor, Statistics, Purdue University

Abstract
Currently, differential privacy (DP) has arisen as the stat-of-the-art method in privacy protection, seeing implementations by tech companies such as Apple, Google, and Facebook, as well as by the U.S. Census. In order to achieve a differential privacy guarantee, an analysis method must introduce additional randomness into the calculations in order to obscure the contributions of any particular individual. While DP methods give a very strong privacy guarantee, the extra randomness makes it challenging to conduct valid statistical inference (such as unbiased estimation, confidence intervals, hypothesis tests, and Bayesian analyses). In this talk, we discuss the challenges of developing valid statistical inference on privatized data and highlight some new solutions to achieve these goals. These results can greatly improve the utility of privatized data, enabling privacy protections in settings that were previously impractical.

About Dr. Awan
Dr. Awan studied at Clarion University from 2011-2014, earning a B.S. in Mathematics. He completed a M.A. in Mathematics at Brandeis University in 2016 under the advisement of Dr. Olivier Bernardi. In May of 2020, He completed his Ph.D. in Statistics at Penn State University, advised by Dr. Aleksandra Slavkovic and Dr. Matthew Reimherr. Currently, Dr. Awan is an Assistant Professor of Statistics at Purdue University. He also works as a differential privacy consultant for the federal non-profit, MITRE.

Lunch Break

12:05p - 1:10p
PMU

Purdue Memorial Union's Atlas Family Marketplace has a wide variety of dining options to satisfy you.

Aatish - Halal contemporary kitchen

BBQ District - Slow-cooked meats, regional sauces and savory sides.

Chef Bill Kim's - Asian dumplings and bowls using authentic ingredients.

Fresh Fare - Fresh flavors with an emphasis on dairy-free and gluten-free options.

Latin Inspired - South American flavors and Latin fare from Brazil and Argentina.

Pizza & Parm Shop - Detroit-style pizza with a caramelized cheese crust and creative parm sandwiches.

Sol Toro - Mexican flavors with a modern flair.

Starbucks®

Sushi Boss - Fresh custom sushi.

Walk On's Sports Bistreaux - Louisiana-inspired cuisine with a game-day flair: burgers, wraps, salads, seafood specialties.

Zen - Build-your-own sushi in a bowl, salad bar and boba teas.

Also... see page 37 for a map of nearby restaurants

Lightning Talk

1:10p
STEW 214

"Assuring High Consequence Systems at Sandia National Labs"

Dr. Ruby E. Booth, Principal Member, Technical Staff, Sandia National Laboratories

Abstract

Our nation's high consequence systems (HCS) are at risk from an ever-increasing digital threat. Digital technologies integrated within HCS pose unique challenges to national security by introducing unexpected behaviors that put missions at risk. These challenges are exacerbated by the growing complexity and interdependence of HCS – characteristics that enable agility and performance in HCS. Sandia's HCS design and development responsibilities are integral to the nation's national security missions, which must be performed reliably even under digital threat.

The United States' national security community recognizes that we lack the technical capabilities needed to make confident, evidence-based assertions of digital assurance efficiently and across high consequence systems. The Digital Assurance for High Consequence Systems Mission Campaign (DAHCS MC, pronounced "Dax") seeks to create this technical basis, and to discover scientific foundations that generalize such capabilities to modern engineering of HCS. With a $45M, seven-year MC, we seek to create the technical basis for efficiently assessing the digital/cyber assurance of the nation's critical systems generally. This includes developing the science-based tools and methods to efficiently 1) characterize, assess, and manage digital risk and 2) design and construct systems with minimal digital risk. Ultimately, our goal is to develop the foundations that allow digital assurance to be incorporated into disciplined systems engineering frameworks.

About Dr. Booth

Dr. Ruby E. Booth is a principal member of the technical staff at Sandia National Laboratories, where she serves as a cybersecurity analyst and national security subject matter expert. She specializes in the interaction between human behavior and cybersecurity. She received her undergraduate degree from Rhodes College and an MS and PhD from the University of Memphis. She is a nonresident fellow of the Berkeley Risk and Security Lab and the Center for Long-Term Cybersecurity at the University of California, Berkeley. including AI/ML, cryptography, trusted hardware, and software engineering. He has published papers in top security, AI, and networking conferences and prestigious journals, such as USENIX Security, NeurIPS, IEEE INFOCOM, IEEE TIFS, IEEE TDSC, IEEE TPDS, and IEEE TSC. He is the recipient of the prestigious NSF CAREER Award. He won the Distinguished Paper Award in ACM ASIACCS 2013. Dr. Sun is a Senior Member of IEEE.

CERIAS Tech Talk

1:30p
STEW 214

"Toward an Internet of Secure Things"

Dr. Saurabh Bagchi, Professor of Electrical and Computer Engineering, Purdue University

Abstract

Cyber-Physical Systems (CPS) are typically composed of interconnected hardware and software components, which individually may not be inherently highly reliable or secure. However, several CPS applications demand a high degree of safety, security, and reliability. Thus, our grand challenge problem is

There has been enormous progress in understanding and patching various classes of vulnerabilities in large-scale distributed CPS. However, these efforts at designing and operating resilient CPS have often been stymied by the lack of understanding of the impact of any failure to the overall system, under the economic and policy constraints involved.

I will first provide the perspective on progress that has been made in securing such large-scale CPS. Then I will present our approach from two lenses. The first lens is a macroscopic one [TCNS-20, AsiaCCS-21, S&P-22] where we look at the security of interdependent CPS managed by multiple defenders that are under the threat of stepping-stone attacks. We model such systems via game-theoretic models and incorporate in them the biases in human decision-making. We present learning techniques for enhancing decision-making in multi-round setups. The second lens is that of stochastic learning [AsiaCCS-23, CVPR-23, S&P-24, CVPR-24] by many CPS devices in a distributed and collaborative manner. We show the vulnerability of such learning, from a security and a privacy standpoint, and then point out some promising defense approaches that can stay within the resource bounds of these devices.

Time permitting, I will discuss a third lens, a microscopic one [UsenixSec-18, NDSS-20, UsenixSec-20] where we look at protecting each individual device against powerful control flow attacks by bringing to low-end devices the core security principle of least privilege execution.

About Dr. Bagchi

Saurabh Bagchi is a Professor in the School of Electrical and Computer Engineering and the Department of Computer Science at Purdue University in West Lafayette, Indiana. His research interest is in dependable computing and distributed systems. He is the founding Director of a university-wide resilience center at Purdue called CRISP (2018-present) and Director of the Army's Artificial Intelligence Innovation Institute (A2I2) (2020-25). He serves on the IEEE Computer Society Board of Governors.

Saurabh is proudest of the 25 PhD students and about 30 Masters thesis students who have graduated from his research group and who are in various stages of building wonderful careers in industry or academia. In his group, he and his students have way too much fun building and breaking real systems. Along the way this has led to 13 best paper awards or runner-up awards at IEEE/ACM conferences and a Test of Time Award. Saurabh serves as the founder and CTO of a cloud computing startup, KeyByte (2022). Saurabh received his MS and PhD degrees from the University of Illinois at Urbana-Champaign and his BS degree from the Indian Institute of Technology Kharagpur, all in Computer Science.

Lightning Talk

2:15p

STEW 214

"Hierarchical Adversarial Inverse Reinforcement Learning"

Dr. Vaneet Aggarwal, Professor, Industrial Engineering, Purdue University

Abstract

In this talk, we will provide novel approaches for Multi-task Imitation Learning, which aims to train a policy capable of performing a distribution of tasks based on multi-task expert demonstrations, which is essential for general-purpose robots. Existing Multi-task imitation learning algorithms suffer from low data efficiency and poor performance on complex long-horizon tasks. We will describe Multi-task Hierarchical Adversarial Inverse Reinforcement Learning, which learns hierarchically-structured multi-task policies, which is more beneficial for compositional tasks with long horizons and has higher expert data efficiency through identifying and transferring reusable basic skills across tasks. Further, this approach leads to generalization to new tasks or transfer learning to similar domains. We will evaluate the performance of the proposed approach using a series of challenging multi-task settings from Mujoco, a widely-utilized environment for simulating and controlling robotic systems.

About Dr. Aggarwal

Vaneet Aggarwal received the BTech degree from the Indian Institute of Technology Kanpur, Kanpur, India, in 2005 and the MA and PhD degrees from Princeton University, Princeton, NJ, USA, in 2007 and 2010, respectively, all in Electrical Engineering. He is currently a Professor in the School of Industrial Engineering, and the School of Electrical and Computer Engineering (by courtesy) at Purdue University, where he has been since Jan 2015.

Prior to joining Purdue, he was a Senior Member of the Technical Staff - Research for five years with AT&T Labs Research, Bedminster, NJ, USA (2010-2014). Dr. Aggarwal has been Adjunct Assistant Professor in the Department of Electrical Engineering at Columbia University (2013-2014), a VAJRA Adjunct Professor in the Department of Electrical Communications Engineering at IISc Bangalore (2018-2019), Visiting Professor at Plaksha University (2022-2023), Adjunct Professor in CS at IIIT Delhi (2022-2023), and Visiting Professor in CS at KAUST (2022-2023). Dr. Aggarwal received the Princeton University's Porter Ogden Jacobus Honorific Fellowship in 2009, AT&T Vice President Excellence Award in 2013, AT&T Senior Vice President Excellence Award in 2014, and Purdue University's Most Impactful Faculty Innovator award in 2020. In addiiton, he received IEEE Jack Neubauer Memorial Award in 2017, IEEE Infocom Workshop Best paper award in 2018, and Neurips Workshop Best paper award in 2021. He was an Associate Editor for the IEEE Transactions on Green Communications and Networking (2017-2020) and IEEE Transactions on Communications (2016-2021). He is currently serving on the Editorial Board of the IEEE/ACM Transactions on Networking (2019-current), and is co-Editor-in-Chief of the ACM Journal on Transportation Systems (2022-current).

Lightning Talk

2:35p
STEW 214

"An Interview Study on Third-Party Cyber Threat Hunting Processes in the U.S. Department of Homeland Security (USENIX Security 2024)"

Dr. James Davis, Assistant Professor
The Elmore Family School of Electrical and Computer Engineering , Purdue University

Abstract

This talk is about cyber threat hunting, which is an emerging discipline of cybersecurity operations. Traditional cybersecurity defense is reactive. Cybersecurity operations centers keep out adversaries and incident response teams clean up after break-ins. Recently a proactive stage has been introduced: Cyber Threat Hunting (TH) looks for potential compromises missed by other cyber defenses. TH is mandated for federal executive agencies and government contractors. As threat hunting is a new cybersecurity discipline, the practices and challenges of TH have not yet been documented. To address this gap, we conducted the first interview study of threat hunt practitioners. We obtained access and interviewed 11 threat hunters associated with the U.S. government's Department of Homeland Security. We describe the diversity among their processes, show that their processes differ from the TH processes reported in the literature, and unify our subjects' descriptions into a single TH process. We enumerate common TH challenges and solutions according to the subjects. The two most common challenges were difficulty in assessing a Threat Hunter's expertise, and developing and maintaining automation. We conclude with recommendations for TH teams (improve planning, focus on automation, and apprentice new members) and highlight directions for future work (finding a TH process that balances flexibility and formalism, and identifying assessments for TH team performance). Our findings will be presented at USENIX Security 2024.

About Dr. Davis

James C. Davis is an assistant professor of Electrical and Computer Engineering at Purdue University and a senior member of the IEEE. He worked for IBM from 2012-2015 and received his PhD degree from Virginia Tech in 2020. His research is published at the most prestigious venues in software engineering (e.g., ICSE, FSE) and cybersecurity (e.g., IEEE S&P, USENIX Security). His work has been recognized with three ACM SIGSOFT distinguished paper awards. His lab is supported by the US National Science Foundation, Google, Rolls Royce, and Cisco.

Networking Break

2:55p
STEW 214

Panel Discussion #3

3:10p

STEW 214

"Where Code Meets Chip"

Moderator - Dr. Anand Raghunathan, Silicon Valley Professor and Chair of the VLSI area in the School of Electrical and Computer Engineering, Purdue University

Anand Raghunathan received the B. Tech. degree in Electrical and Electronics Engineering from the Indian Institute of Technology, Madras, India, and the M.A. and Ph.D. degrees in Electrical Engineering from Princeton University, Princeton, NJ. He is currently the Silicon Valley Professor and Chair of the VLSI area in the School of Electrical and Computer Engineering at Purdue University. He serves as Associate Director of the $36M SRC/DARPA Center for Brain-inspired Computing (C-BRIC) and founding co-director of the Purdue/TSMC Center for a Secured Microelectronics Ecosystem (CSME). His research explores brain-inspired computing, energy-efficient and high-performance machine learning, system-on-chip design and computing with post-CMOS devices. He holds a Distinguished Visiting Chair at the Indian Institute of Technology, Madras, where he is helping establish a Center for Computational Brain Research.

Max Dewees
Product Security Architect, Analog Devices, Inc.

Max DeWees is a Product Security Architect at Analog Devices, Inc. He has been with ADI for 10 years, after receiving his master's degree in Information Security from Purdue University. Before his current role, Max led an embedded firmware product development team, creating device drivers and libraries for cryptographic hardware accelerators and secure bootloaders for silicon products in the automotive, industrial, and consumer markets. Max's current focus is helping to define next-generation security architecture for ADI's upcoming Intelligent Edge platforms.

Dr. Nathaniel Husted, Chief Scientist Cyber & EW technology, Naval Surface Warfare Center, CRANE

Dr. Nathaniel Husted serves as Chief Scientist for Cyber and Electromagnetic Warfare (EW) Technologies in the Expeditionary EW division at Naval Surface Warfare Center – Crane in southern Indiana. His duties focus on growing Crane's workforce in the area of Cyber/EW, developing infrastructure to enable left-of-acquisition Cyber/EW activities, and helping bridge the gap between the Cyberspace and EMSO communities. Prior to his current role as Chief Scientist he served as the national Cybersecurity Engineering lead for a major new Navy program and as the first program officer (acting) for Expeditionary Cyber at the Office of Naval Research (ONR). He obtained a B.S. from Purdue University at the Indianapolis campus and a Ph.D. in Informatics from Indiana University. His research focused on the intersection of Computer Security and Complex Systems with further contributions in the areas of applied cryptography and the macroeconomics of information security. He has contributed to the Linux Kernel and Buildroot open-source projects. His most recent internally funded project culminated in a hardware-in-the-loop, high assurance, flight system for teaching Crane's workforce cybersecurity and formal assurance concepts. The project leverages Kerbal Space Program and the Rust programming language; it is available at: https://github.com/NSWC-Crane/kerbx-flightsystems. He currently has returned to ONR to support the Expeditionary Cyber portfolio as a strategic technical advisor.

Dr. Dave Tian
Assistant Professor
Computer Science, Purdue University

Dave (Jing) Tian is an Assistant Professor in the Department of Computer Science at Purdue University working on systems security and a faculty at the PurSec Lab,
where he and his colleagues supervise over 30 Ph.D. students covering all layers of systems security.
His research involves embedded systems security, operating systems security, trusted and confidential computing, and hardware security and trust.
He has published 50+ peer-reviewed conference papers and journals,
including over 25 top-tier security conference papers (IEEE S&P, USENIX Security, ACM CCS, ISOC NDSS).
His research is funded by ONR, DARPA, NSF, Intel, Lockheed Martin, Rolls-Royce, etc.
He received an NSF CAREER award in 2022.

Lightning Talk

4:10p

STEW 214

"Packing Arithmetic to Improve Security and Performance of Privacy Technologies"

Dr. Hemanta Maji, Associate Professor of Computer Science, Purdue University

Abstract
Use cases in privacy like identification, authentication, private transactions, and verifiable/secure computation rely on cryptography to handle secrets. However, incompatibility of the application front-end and cryptographic back-end operations leads to efficiency and security bottlenecks.

1. For instance, in the zero-knowledge (ZK) domain, Boolean circuits generate auxiliary information like hashes and Merkle trees in privacy-preserving cryptography, blockchain systems, and outsourced computation. However, ZK back-ends support arithmetic only over large fields -- not Boolean arithmetic.
2. Likewise, in secure computation, it is challenging to evaluate Boolean circuits securely using homomorphic cryptosystems like Ring-LWE and Paillier-type schemes that only support large arithmetic.

Arithmetic packing, our information-theoretic invention, resolves these incompatibilities. It is a versatile intermediate representation that implements vector instructions for application front-ends using a single large back-end arithmetic. This invention has led to faster and more secure MPC and ZK technologies.

This talk will introduce arithmetic packings, their applications, our new combinatorial and algebraic complexity-theoretic framework to investigate these new packing problems, and optimal packing constructions relying on computer-assisted search.

About Dr. Maji
Maji is an Associate Professor in the Department of Computer Science at Purdue University. He is interested in Cryptography and Algorithms, particularly Secure Computation and Information-theoretic Cryptography. Among others, his current research makes zero-knowledge and secure computation technologies faster and safer. of the IEEE. He worked for IBM from 2012-2015 and received his PhD degree from Virginia Tech in 2020. His research is published at the most prestigious venues in software engineering (e.g., ICSE, FSE) and cybersecurity (e.g., IEEE S&P, USENIX Security). His work has been recognized with three ACM SIGSOFT distinguished paper awards. His lab is supported by the US National Science Foundation, Google, Rolls Royce, and Cisco.

Closing Keynote

4:30p

STEW 214

"Where Code Meets Chip"

Philippe Biondi is an executive expert in cyber security at Airbus. He is a co-creator of the SSTIC french-speaking conference. He is the author of Scapy and numerous other security related tools. He is a co-author of the Security Powertools book.
He published several articles in MISC magazine. He gave several talks to security conferences (Blackhat, HITB, GreHack, CansecWest, Defcon, Syscan, etc.).

Poster Session Abstracts

POSTERS

ASSURED IDENTITY AND PRIVACY 38

END SYSTEM SECURITY 41

HUMAN CENTRIC SECURITY 43

POSTER SESSION RESEARCH AREA KEY

Artificial Intelligence	Red
Assured Identity and Privacy	Blue
End System Security	Pink
Human Centric Security	Yellow
Network Security	Violet
Prevention, Detection and Response	Green
Policy, Law and Management	Gold

These posters, and posters from previous years, are available at https://ceri.as/posters

ARTIFICIAL INTELLIGENCE

1. Achieving Algorithmic Fairness through Label Flipping

Shashank Thandri

As machine learning (ML) and artificial intelligence (AI) become increasingly prevalent in high-stake decision making, fairness has emerged as a critical societal issue. Individuals belonging to diverse groups receive different algorithmic outcomes largely due to the inherent errors and biases in the underlying training data, thus resulting in violations of group fairness or bias. We address the problem of resolving group fairness by flipping the labels of instances in the training data. We propose solutions to obtain an ordering in which the labels of training data instances should be flipped to reduce the bias in predictions of a model trained over the modified data. We experimentally evaluate our solutions on several real-world datasets and demonstrate that bias is reduced by flipping a small fraction of training data labels.

2. An Interactive Framework for Profiling News Media Sources

Nikhil Mehta and Dr. Dan Goldwasser
mehta52@purdue.edu, dgoldwas@purdue.edu

The recent rise of social media has led to the spread of large amounts of fake and biased news, content published with the intent to sway beliefs. While detecting and profiling the sources that spread this news is important to maintain a healthy society, it is challenging for automated systems. In this paper, we propose an interactive framework for news media profiling. It combines the strengths of graph based news media profiling models, Pre-trained Large Language Models, and human insight to characterize the social context on social media. Experimental results show that with as little as 5 human interactions, our framework can rapidly detect fake and biased news, even in the most challenging settings of emerging news events, where test data is unseen.

3. Data Acquisition to Improve Machine Learning Fairness through Multi-Armed Bandit

Jahid Hasan

Over the past few decades, the extensive use of machine learning (ML) has shifted our focus from its implementation to its consequences. Several instances have indicated bias in ML-based systems deployed in sensitive fields (e.g., law, finance, HR, etc.), which has become a matter of concern. Numerous studies have shown that these biases in ML models originate from biased training data, making data the root cause of the issue. Several existing data preparation studies can address this issue. However, these approaches are problem-specific and can negatively impact downstream data usage. A more efficient approach would be to focus on earlier stages in the data science pipeline, such as data acquisition, which can significantly improve the quality of downstream analyses. To address this, we employ a comprehensive solution for fair data acquisition that includes data source selection, merging sources, clustering data instances, and finally, adopting an approach based on multi-armed bandits to acquire data for improved model fairness.

4. Discovering Adversarial Driving Maneuvers against Autonomous Vehicles

Ruoyu Song, Muslum Ozgur Ozmen, Hyungsub Kim, Raymond Muller, Z. Berkay Celik, and Dr. Antonio Bianchi

Over 33% of vehicles sold in 2021 had integrated autonomous driving (AD) systems. While many adversarial machine learning attacks have been studied against these systems, they all require an adversary to perform specific (and often unrealistic) actions, such as carefully modifying traffic signs or projecting malicious images, which may arouse suspicion if discovered. In this paper, we present Acero, a robustness-guided framework to discover adversarial maneuver attacks against autonomous vehicles (AVs). These maneuvers look innocent to the outside observer but force the victim vehicle to violate safety rules for AVs, causing physical consequences, e.g., crashing with pedestrians and other vehicles. To optimally find adversarial driving maneuvers, we formalize seven safety requirements for AD systems and use this formalization to guide our search. We also formalize seven physical constraints that ensure the adversary does not place themselves in danger or violate traffic laws while conducting the attack. Acero then leverages trajectory-similarity metrics to cluster successful attacks into unique groups, enabling AD developers to analyze the root cause of attacks and mitigate them. We evaluated Acero on two open-source AD software, openpilot and Autoware, running on the CARLA simulator. Acero discovered 219 attacks against openpilot and 122 attacks against Autoware. 73.3% of these attacks cause the victim to collide with a third-party vehicle, pedestrian, or static object.

5. Estimating Machine Learning Model Fairness through Data Characteristics

Kevin Chittilapilly, Ahana Bhattacharya

The pursuit of fairness in machine learning (ML) systems is a critical challenge in todays world that relies heavily on AI systems. However computing the fairness necessitates substantial computational resources and time when evaluating across entire datasets. This research introduces an innovative approach to estimate fairness in ML systems by leveraging data characteristics and constructing a metafeatures dataframe. Using our methodology enables the prediction of fairness with significantly reduced computational cost and expedited analysis times. Furthermore, we explore the application of data models as an alternative to traditional machine learning techniques for predicting fairness. This dual approach not only enhances the efficiency of fairness assessments in ML systems but also provides a scalable framework for future fairness evaluation methodologies. Our findings suggest that using data characteristics to estimate fairness is not only feasible but also effective, offering a promising avenue for developing more equitable ML systems with reduced resource consumption.

6. Eureka: A General Framework for Black-box Differential Privacy Estimators

Yu Wei, Dr. Vassilis Zikas

Differential privacy (DP) is a key tool in privacy-preserving data analysis. Yet it remains challenging for non-privacy-experts to prove the DP of their algorithms. We propose a methodology for domain experts with limited data privacy backgrounds to empirically estimate the privacy of an arbitrary mechanism. Our Eureka moment is a new link---which we prove---between the problems of DP parameter-estimation and Bayes optimal classifiers in ML, which we believe can be of independent interest. Our estimator uses this link to achieve two desirable properties: (1) black-box, i.e., it does not require knowledge of the underlying mechanism, and (2) it has a theoretically-proven accuracy, depending on the underlying classifier used, allowing plug-and-play use of different classifiers. More concretely, motivated by the impossibility of the above task for unrestricted input domains (which we prove), we introduce a natural, application-inspired relaxation of DP which we term relative DP. Intuitively, relative DP defines a mechanism's privacy relative to an input set T, circumventing the above impossibility when T is finite. Importantly, it preserves the key intuitive privacy guarantee of DP while enjoying a number of desirable DP properties---scalability, composition, and robustness to post-processing. We then devise a black-box poly-time (ϵ,δ)-relative DP estimator for any poly-size T---the first privacy estimator to support mechanisms with large {\em output} spaces while having tight accuracy bounds. As a result of independent interest, we generalize our theory to develop the first Distributional Differential Privacy (DDP) estimator. We benchmark our estimator in a proof-of-concept implementation. First, using kNN as the classifier we show that our method (1) produces a tight, analytically computed (ϵ, δ)-DP trade-off of low-dimensional Laplace and Gaussian mechanisms---the first to do so, (2) accurately estimates the privacy spectrum of DDP mechanisms, and (3) can verify a DP mechanism's implementations, e.g., Sparse Vector Technique, Noisy Histogram, and Noisy max. Our implementation and experiments demonstrate the potential of our framework, and highlight its computational bottlenecks in estimating DP, e.g., in terms of the size of δ and the data dimensionality. Our second, neural-network-based instantiation makes a first step in showing that our method can be extended to mechanisms with high-dimensional outputs.

7. Generative AI and Open-Source Intelligence: Exploring Capabilities and Privacy Implications

Akif Ozer, Vinicius Lima, Dr. Umit Karabiyik

In the rapidly evolving landscape of artificial intelligence, the emergence and growing accessibility of generative AI technologies have marked a significant milestone. Tools like Gemini and ChatGPT 4, which can search the internet and harness online information, have opened new avenues for research and application in various fields. This study explores generative AI technologies' open-source intelligence (OSINT) capabilities, highlighting their potential in gathering and analyzing information from publicly accessible digital spaces. Despite these undeniable advantages, these tools possess built-in privacy protections to mitigate the risks of accessing and disseminating personal information. Our research methodology involves rigorous testing of these tools to assess their effectiveness in OSINT tasks while evaluating the robustness of their privacy protection mechanisms. Through this study, we aim to provide insights into the balance between the powerful capabilities of generative AI technologies and the ethical considerations of privacy and data protection. The findings of this research contribute to the broader understanding of the role of generative AI in the future of information gathering and analysis and stimulate discussion on the responsible development and use of these technologies.

8. LeMix: Rehosting Embedded Systems as Linux Application for Effective Vulnerability Detection

Sai Ritvik Tanksalkar, Jayashree Srinivasan, Srihari Danduri, Paschal C. Amusuo, James C. Davis, Dr. Aravind Machiry

Dynamic analysis stands out as a crucial capability for security assessment in embedded systems software. Most dynamic analysis techniques focus on emulation. However, this approach encounters scalability challenges due to the diversity of embedded hardware and software platforms. We introduce Lemix, a framework to separate an embedded application from its hardware dependencies, allowing it to operate on a standard Linux platform. Lemix makes it easier to find and fix security issues in the application using traditional security analysis techniques. Our key insight is that targeted modifications to an application's source code, aimed at eliminating architecture-dependent dependencies, generally do not significantly impact the application's functionality or the underlying layers of the operating system it relies on. Our results further validate that maintaining high application fidelity is not necessarily a prerequisite for effective security analysis.

9. LOKI: Large-scale Data Reconstruction Attack against Federated Learning

Joshua C. Zhao, Ahaan Dabholkar, Atul Sharma

Federated learning was introduced to enable machine learning over large decentralized datasets while promising privacy by eliminating the need for data sharing. Despite this, prior work has shown that shared gradients often contain private information and attackers can gain knowledge either through malicious modification of the architecture and parameters or by using optimization to approximate user data from the shared gradients. However, prior data reconstruction attacks have been limited in setting and scale, as most works target FedSGD and limit the attack to single-client gradients. Many of these attacks fail in the more practical setting of FedAVG or if updates are aggregated together using secure aggregation. Data reconstruction becomes significantly more difficult, resulting in limited attack scale and/or decreased reconstruction quality. When both FedAVG and secure aggregation are used, there is no current method that is able to attack multiple clients concurrently in a federated learning setting. In this work we introduce LOKI, an attack that overcomes previous limitations and also breaks the anonymity of aggregation as the leaked data is identifiable and directly tied back to the clients they come from. Our design sends clients customized convolutional parameters, and the weight gradients of data points between clients remain separate even through aggregation. With FedAVG and aggregation across 100 clients, prior work can leak less than 1% of images on MNIST, CIFAR-100, and Tiny ImageNet. Using only a single training round, \name is able to leak 76-86% of all data samples.

10. Malware Language Processing "MLP": Developing a new paradigm for malware analysis and classification using Machine Learning and Artificial Intelligence

Solomon Sonya

Malware continues to increase in prevalence and sophistication. Successfully exploiting networks and digital systems has become a highly profitable operation for malicious threat actors. VirusTotal reported a daily submission of 2M+ malware samples in March 2024 (VirusTotal, 2024). Of those 2 million daily submissions, over 1 million were unique malware samples (per day!). Traditional detection mechanisms including antivirus software fail to adequately detect new and varied malware (Jhaveri et al, 2022, Johnson and Haddad, 2021, Geis, 2019). Artificial Intelligence and Machine Learning models provide advanced capabilities that can enhance cybersecurity. Building a robust and automated artificially intelligent malware analysis pipeline and producing new malware datasets however, are not trivial. The purpose of this poster is to present current progress in our research aimed at developing a robust malware classification framework. We are developing this framework to automate malware analysis and feature extraction and produce new, standardized malware datasets for future Machine Learning analysis. Additionally, this research presents status regarding the development of a new Malware Ensemble Classification Facility that leverages several Machine Learning models to enhance the classification of malware. To our knowledge, this is the first research that utilizes Machine Learning to provide enhanced classification of an entire 200+ gigabyte, malware family corpus consisting of 80K+ unique malware samples and 70+ malware families

11. Mechanism Design for Control-theoretic Objectives

Dr. Vijay Gupta, Mostafa M. Shibl

The creation of local control laws for each individual agent in a multiagent system is crucial to ensuring that the emergent global behavior is desired in relation to a specific system level aim. Specifically, we derive a methodology for creating local agent objective functions that ensures that the resulting game has an inherent structure that can be leveraged in distributed learning applications, such as Markov potential games, and that there is an equivalence between the optimizers of the system level objective and the resulting equilibria. This allows any distributed learning approach that ensures convergence to an equilibrium for the obtained game structure to be used to finish the control design. Thus, the problem aims to leverage multi-agent reinforcement learning algorithms such as policy gradient algorithms to control dynamical systems through game theoretic approaches.

12. Modeling and Detecting Falsified Vehicle Trajectories Under Data Spoofing Attacks

Jun Ying, Dr. Yiheng Feng, Qi Alfred Chen, Z. Morley Mao

Connected Vehicle (CV) and Connected and Autonomous Vehicle (CAV) technologies can greatly improve traffic efficiency and safety. Data spoofing attack is one major threat to CVs and CAVs, since abnormal data (e.g., falsified trajectories) may influence vehicle navigation and deteriorate CAV/CV-based applications. In this work, we aim to design a generic anomaly detection model which can be used to identify abnormal trajectories from both known and unknown data spoofing attacks. First, the attack behaviors of two representative and sophisticated known attacks are modeled. Then, Using driving features derived from transportation and vehicle domain knowledge, an anomaly detection framework is proposed. The framework combines a feature extractor and an anomaly classifier trained with known attack trajectories and can be applied to identify falsified trajectories generated by various attacks. In the numerical experiment, a highway segment with a signalized intersection is built in the V2X Application Spoofing Platform (VASP). To evaluate the generality of the proposed anomaly detection algorithm, we further tested the proposed model with several unknown attacks provided in VASP. The results indicate that the proposed model achieves high accuracy in detecting falsified attack trajectories from both known and unknown attacks and outperforms several baselines. Furthermore, we show the importance of integrating domain knowledge in the feature selection process.

13. Nanomanufactured connected wearable sensors for human body digital twins

Dr. Wenzhuo Wu, Jing Jiang

Our group focuses on innovating unprecedented material technologies for next-generation products through nanomanufacturing across length scales. Our core expertise lies in nanomanufacturing, semiconductor fabrication, and data-driven design. By leveraging these capabilities, we aim to develop cutting-edge materials and devices with enhanced performance and functionality. The application areas of our research includes wearable devices, ubiquitous sensors, robotics, quantum electronics, etc. Through interdisciplinary collaboration and advanced manufacturing techniques, we strive to address the pressing challenges in these fields and pave the way for human body digital twins technology.

14. Preserving Fairness Generalization in Deepfake Detection

Li Lin, Xinan He, Yan Ju, Xin Wang, Feng Ding, Dr. Shu Hu

Although effective deepfake detection models have been developed in recent years, recent studies have revealed that these models can result in unfair performance disparities among demographic groups, such as race and gender. This can lead to particular groups facing unfair targeting or exclusion from detection, potentially allowing misclassified deepfakes to manipulate public opinion and undermine trust in the model. The existing method for addressing this problem is providing a fair loss function. It shows good fairness performance for intra-domain evaluation but does not maintain fairness for cross-domain testing. This highlights the significance of fairness generalization in the fight against deepfakes. In this work, we propose the first method to address the fairness generalization problem in deepfake detection by simultaneously considering features, loss, and optimization aspects. Our method employs disentanglement learning to extract demographic and domain agnostic forgery features, fusing them to encourage fair learning across a flattened loss landscape. Extensive experiments on prominent deepfake datasets demonstrate our method's effectiveness, surpassing state-of-the-art approaches in preserving fairness during cross-domain deepfake detection.

15. Securing Deep Neural Networks on Edge from Membership Inference Attacks Using Trusted Execution Environments

Cheng-yun Yang

Privacy concerns arise from malicious attacks on Deep Neural Network (DNN) applications during sensitive data inference on edge devices. Membership Inference Attack (MIA) is developed by adversaries to determine whether sensitive data is used to train the DNN applications. Prior work uses Trusted Execution Environments (TEEs) to hide DNN model inference from adversaries on edge devices. Unfortunately, existing methods have two major problems. First, due to the restricted memory of TEEs, prior work cannot secure large-size DNNs from gradient-based MIAs. Second, prior work is ineffective on output-based MIAs. To mitigate the problems, we present a depth-wise layer partitioning method to run large sensitive layers inside TEEs. We further propose a model quantization strategy to improve the defense capability of DNNs against output-based MIAs and accelerate the computation. We also automate the process of securing PyTorch-based DNN models inside TEEs. Experiments on Raspberry Pi 3B+ show that our method can reduce the accuracy of gradient-based MIAs on AlexNet, VGG-16, and ResNet-20 evaluated on the CIFAR-100 dataset by 17.9%, 11%, and 35.3%. The accuracy of output-based MIAs on the three models is also reduced by 18.5%, 13.4%, and 29.6%, respectively.

16. The Trustworthiness of Large Language Models in Long Context Recall

Yifei Hu

Long Context QA has become one of the common use cases for Generative Large Language Models (LLMs). We challenge the LLMs with a practical long context recall and reasoning task by asking the LLMs to repeat certain sentences from a given long document. Our results suggest that even the state-of-the-art LLMs still cannot perfectly recall the original text in this setting.

17. Valuation-based Data Acquisition to Improve Machine Learning Fairness

Ekta, Dr. Romila Pradhan

Machine learning algorithms are increasingly being used in a variety of applications and are heavily relied upon to make decisions that impact people's lives. ML models are often praised for their precision, yet they can discriminate against certain groups due to biased data. Historical inequities can propagate through machine learning, posing a challenge to developing models that are fair and unbiased for all. One of the major factors that lead to bias is the data used to train them. It is important to address the biases in the training data, as they can lead to unfair and unjust results when the model is deployed in real-world applications. The induced bias due to data can be mitigated using three methodologies i.e., pre-processing, in-processing, and post-processing. This study investigates Data Acquisition as a potential bias mitigation technique, which is closest to pre-processing in the machine learning pipeline.

18. Vigilante Defender: A Vaccination-based Defense Against Backdoor Attacks on 3D Point Clouds Using Particle Swarm Optimization

Agnideven Palanisamy Sundar, Dr. Feng Li, Dr. Xukai Zou, Dr. Tianchong Gao, Dr. Yucheng Xie, Ryan Hosler.

Backdoor attacks, where hidden triggers in training data cause incorrect model predictions, pose significant threats to point cloud (PC) applications. These stealthy attacks, involving subtle point cloud manipulations, compromise models especially in distributed learning environments where data is pooled for central training. To counteract this, we introduce a novel 'vaccination' strategy that allows data contributors with only black-box model access to independently defend against such attacks. By embedding benign 'vaccination' triggers discovered through Particle Swarm Optimization, our method effectively neutralizes potential backdoors, preserving both the model's integrity and task performance. Tested on standard datasets with common PC models like PointNet and DGCNN, our experiments show a marked decrease in attack success rates with negligible impact on accuracy.

19. XAI-ADS: An Explainable Artificial Intelligence Framework for Enhancing Anomaly Detection in Autonomous Driving Systems

Sazid Nazat, Lingxi Li, PI: Dr. Mustafa Abdallah

The advent of autonomous driving systems has given rise to pressing cybersecurity issues regarding the vulnerability of vehicular ad hoc networks (VANETs) to potential attacks. This critical security problem necessitates the application of artificial intelligence (AI) models for anomaly detection in VANETs of autonomous vehicles (AVs). However, the lack of explainability of such AI-based anomaly detection models presents challenges. This motivates an emerging research direction of utilizing explainable AI (XAI) techniques to elucidate the behaviors of anomaly detection models in AV networks. In this work, we propose an end-to-end XAI framework to interpret and visualize the anomaly detection classifications made by AI models securing VANETs. We evaluate the framework on two real-world autonomous driving datasets. The framework furnishes both global and local explanations for the black-box AI models using two XAI methods. Moreover, we introduce two novel feature selection techniques to identify the salient features contributing to anomaly detection, derived from the popular SHAP XAI method and the accuracy of six different black-box AI models. We compare our proposed feature selection approaches with six state-of-the-art feature selection techniques (including two wrapper-based feature selection methods), demonstrating superior performance on various evaluation metrics. To generalize the impact of our feature selection methods, we apply three independent classifiers to evaluate our proposed feature selection approaches. The novel feature selection methods effectively distill the most explanatory features, enhancing model interpretability. Finally, we assess the efficiency (how quickly the XAI models can yield explanatory findings) for each of the six black-box AI models we employed on our two datasets, identifying the most efficient model. By furnishing explanations and visualizations of anomaly detection by AI models, our XAI framework can help in enabling trust and transparency for securing vehicular networks.

ASSURED IDENTITY AND PRIVACY

20. A Policy-Agnostic Language for Oblivious Computation

Qianchuan Ye, Dr. Benjamin Delaware

Secure multiparty computation (MPC) techniques allow multiple parties to collaboratively compute functions over sensitive data in a privacy-preserving manner. MPC protocols use powerful cryptographic techniques to achieve these privacy guarantees, making them challenging for non-experts to directly use. To address this challenge, several high-level languages have been proposed to make writing such applications accessible. These languages typically require the programmers to embed their privacy policies into the application logic, making it hard to audit the policies, or experiment with different policies. This poster presents our ongoing development of a privacy-preserving language, Taype, that decouples privacy and functionality concerns. Two key ingredients of this language are oblivious algebraic data types and tape semantics. Oblivious algebraic data types are a form of dependent types with oblivious constructs, that can be used to modularly encode complex privacy policies for structured data. Tape semantics then enforce these policies during execution, enabling applications to modularly compose policies and programs written in a conventional way without compromising privacy.

21. Directed Infusion of Data (DIOD) for Secure Data Transfer

Tyler Lewis, Dr. Arvind Sundaram

The emergence of AI/ML tools heavily incentivizes collaboration among industrial partners and research institutions to ensure that the vast quantities of data they possess are efficiently leveraged. Large dynamic systems such as power plants, grid applications, smart factories, etc., require extensive engineering analytics such as condition monitoring, time series prediction, control implementation, etc. However, the possibility of data leakage due to a malicious third-party or an untrustworthy collaborator leads to major security concerns that may prevent collaboration. While most collaborators are assumedly trustworthy, it is difficult to absolutely ensure that data cannot be misused; even when the analyst is trustworthy, sensitive data passed to them is vulnerable to their network's security. While various researchers have experimented with techniques to bypass this assumption, the current state-of-the-art methods inherently limit the results of collaboration, inducing a tradeoff between data privacy and fidelity of results. DIOD allows efficient data masking that cannot be reverse engineered by any third-party, regardless of their knowledge of the system. In addition, the data's inferential properties are preserved by DIOD, permitting complex analyses, (e.g., classification, regression, etc.), to be performed on the masked data directly. DIOD has been studied across a variety of circumstances, including condition monitoring tasks, simple binary classification problems, and linear regression.

22. More is Merrier: Relax the Non-Collusion Assumption in Multi-Server PIR

Tiantian Gong, Ryan Henry, Alexandros Psomas, Dr. Aniket Kate

A long line of research on secure computation shows that anything that can be computed, can be computed securely using a set of non-colluding parties. This non-collusion assumption is pervasive across secure multi-party computation (MPC). But it remains highly susceptible to covert, undetectable collusion among computing parties. In this work, we relax this traditional assumption in the context of multi-server 1-private PIR (Private Information Retrieval), a crucial aspect of privacy-preserving computations. Traditionally, it operates under the assumption of no pair-wise collusion. Our work introduces a novel collusion deterrence mechanism designed, analyzed, and implemented specifically for 1-private PIR on a public bulletin board, meticulously considering rational and malicious parties.

23. Semi Differential Privacy

Young Hyun Cho, Dr. Jordan Awan

Differential privacy (DP) is the state-of-the-art framework for formal privacy protection and has become the gold standard for rigorous privacy guarantees. However, it is not always implemented exactly especially when the exact statistics is published along with the DP mechanism. For example, US Census publishes a combination of DP and exact statistics. There is currently no satisfactory method of quantifying the privacy in such settings. In this regard, we propose a framework semi-DP, which properly accounts for the combination of private and non-private releases. Under semi-DP we derive optimal mechanisms under a variety of scenarios.

24. zk-creds: Flexible Anonymous Credentials from zkSNARKs and Existing Identity Infrastructure

Michael Rosenberg, Jacob White, Dr. Christina Garman, Ian Miers

Frequently, users on the web need to show that they are, for example, not a robot, old enough to access an age restricted video, or eligible to download an ebook from their local public library without being tracked. Anonymous credentials were developed to address such concerns. However, existing schemes do not handle the realities of deployment or the complexities of real-world identity assurance. Instead, they implicitly make assumptions such as there being an issuing authority for anonymous credentials that, for real applications, requires sophisticated cryptographic tokens. In reality, there are often multiple trust sources for a given identity attribute, their credentials have distinctively different formats, and many, if not all, issuers are unwilling to adopt new protocols. We present and build zk-creds, a protocol that uses general-purpose zero-knowledge proofs to 1) remove the need for credential issuers to hold signing keys: credentials can be issued to a bulletin board instantiated as a transparency log, Byzantine system, or even a blockchain; 2) convert existing identity documents into anonymous credentials without modifying documents or coordinating with their issuing authority; 3) allow for flexible, composable, and complex identity statements over multiple credentials. Concretely, identity assertions using zk-creds take less than 150ms in a real-world scenario of using a passport to anonymously access age-restricted videos.

END SYSTEM SECURITY

25. Ransomware vs Malware Classification Using Subgraph Mining of Function Call Graph

Garvit Agarwal, Dr. Feng Li

In our research, we delve into the often-overlooked difference between ransomware and general malware. It's an important distinction that matters a great deal when it comes to how we defend our digital spaces. By utilizing the keen observational environment of Cuckoo Sandbox, we pull out the patterns of API calls that give away the software's intent. These patterns are mapped out into graphs, each labeled by the type of activity they represent, whether it's a network signal or a file alteration. We then zoom in on the most telling parts of these graphs, converting them into a numerical form that can tell us the story at a glance. To make sense of this data, we crafted a 1D Convolutional Neural Network that learns to tell apart the villains from the mere tricksters. Our work not only sheds light on a nuanced aspect of cybersecurity but also equips systems with sharper tools to prioritize threats and protect our digital environment more effectively.

26. Secure High-Performance Interrupts For Secure High-Performance Processors

Berk Aydogmus

The advent of User level Interrupts sparked an advance in the workloads that benefit from low latency notification systems. User level schedulers, and high throughput devices are the two main customers of such systems. The lack of OS management in interruption, can however cause security exploits. The approach we propose for User Interrupt handling eliminates the latency side-channel identified in prior works, and reduces the overhead of User Interrupts by a significant margin. Moreover, we investigate the potential integration of hardware timers to further enhance user-level preemptive schedulers. While current user level schedulers dedicate a core to generate interrupts on dedicated time intervals, a hardware timer generates interrupts in core, which reduces the constant overhead of interrupt communication bookkeeping, and the cycles spent spinning for the next interrupt interval.

27. Securing Contrastive mmWave-based Human Activity Recognition against Adversarial Label Flipping

Amit Singha, Ziqian Bi, Dr. Tao Li, Yimin Chen, Yanchao Zhang

Wireless Human Activity Recognition (HAR), leveraging their non- intrusive nature, has the potential to revolutionize various sectors, including healthcare, virtual reality, and surveillance. The advent of millimeter wave (mmWave) technology has significantly enhanced the capabilities of wireless HAR systems. This paper presents the first systematic study on the vulnerabilities of mmWave-based HAR to label flipping poisoning attacks in the context of supervised contrastive learning. We identify three label poisoning attacks on the contrastive mmWave-based HAR and propose corresponding countermeasures. The efficacy of the attacks and also our coun- termeasures are experimentally validated on a prototype system. The attacks and countermeasures can be easily extended to other wireless HAR systems, thereby promoting security considerations in system design and deployment.

28. SoK: A Defense-Oriented Evaluation of Software Supply Chain Security

Eman Abu Ishgair, Marcela S. Melara, Dr. Santiago Torres Arias

The software supply chain comprises a highly complex set of operations, processes, tools, institutions and human factors involved in creating a piece of software. A number of high-profile attacks that exploit a weakness in this complex ecosystem have spurred research in identifying classes of supply chain attacks. Yet, practitioners often lack the necessary information to understand their security posture and implement suitable defenses against these attacks. We argue that the next stage of software supply chain security research and development will benefit greatly from a defense-oriented approach that focuses on holistic bottom-up solutions. To this end, this paper introduces the AStRA model, a framework for representing fundamental software supply chain elements and their causal relationships. Using this model, we identify software supply chain security objectives that are needed to mitigate common attacks, and systematize knowledge on recent and well-established security techniques for their ability to meet these objectives. We validate our model against prior attacks and taxonomies. Finally, we identify emergent research gaps and propose opportunities to develop novel software development tools and systems that are secure-by-design.

HUMAN CENTRIC SECURITY

29. Digital Guardian: Harnessing AI to Devalue False Information and Protect Public Discourse on Social Media

Nicholas Harrell

Social media, though used for sharing stories, ideas, and trends, has become a battleground for nefarious entities manipulating public discourse on critical issues. This manipulation, often fueled by viral false information, erodes trust in government, breeds conspiracies, and disrupts critical events like elections and emergency response. Existing research focuses heavily on the emotional appeal of viral content, neglecting the crucial roles of initial discovery context and value alignment within target audiences. This project bridges these gaps by leveraging state-of-the-art artificial intelligence tools to measure online values and their alignment with shared content. By predicting the virality potential based on value alignment, we empower security analyst and policymakers to 1) detect and mitigate foreign disinformation campaigns early, 2) proactively counter harmful narratives with targeted communication, 3) design interventions based on specific community value profiles, and 4) develop public education campaigns promoting critical online thinking. This research holds immense value for policymakers as it equips them with the tools to safeguard national security in the face of evolving online threats.

30. Navigating Software Supply Chain Risks: Practitioner Perspectives on Software Signing

Kelechi G. Kalu, James C. Davis

In today's interconnected software landscape, safeguarding the integrity of software supply chains has become imperative. Software Engineers have witnessed a concerning surge in software supply chain attacks despite the development of various methods, tools, standards, and guidelines to mitigate these attacks. Current literature, regulations, and security frameworks recommend some security baselines, chief amongst them is Software Signing. While software supply chain attack incidences continue to increase, current literature on the software supply chain focuses on both the identification of risk factors and potential attack vectors in Open-source software artifacts, as well as the development of security methods to mitigate these risks. However, we lack an understanding of how Software Supply Chain risks are perceived by practitioners and how proposed security methods, such as Software Signing contribute and are implemented to mitigate them. This knowledge gap poses significant concerns since it may result in inadequate risk management strategies as proposed by regulations (and standards)and inadequate design considerations of these security methods and tools. This potentially exposes software ecosystems and organizations to serious vulnerabilities and security threats. This study conducts a qualitative analysis of software supply chain risks as perceived by practitioners. Additionally, we also investigate the contribution and importance of Software Supply Chain methods contribute to mitigating these risks. Specifically, we conduct a case study on Software Signing as a widely recommended Software Software Supply Chain method. We conducted interviews with 18 practitioners representing 11 organizations to understand how Software Signing in mitigating Software Supply Chain attacks. From our data, we identify Software Supply Chain risks highlighted by practitioners, the importance attached to Software Supply Chain security methods like Software Signing in the Software Engineering Process of various teams, and reasons why certain Software Signing implementations are preferred over others by practitioners. Our findings aim to contribute to the understanding of software supply chain security by highlighting the impact of human factors on Software Supply Chain risks and providing nuanced insights for effectively implementing Software Supply Chain security methods in practice.

31. Perceptions of Cyber Security Student Preparedness

Apoorva Shrivastava, Dr. Tatiana Ringenberg

While students work hard through their undergraduate careers, many still struggle to get a job right after graduation. This study seeks to figure out why and help undergraduate students go the extra mile to help them secure a job. A survey was sent out to industry professionals to help understand their expectations of students and their perspectives on the matter.

32. Personality Traits and Resistance to Online Trust Exploitation

Vaishnavi Mahindra, Tatiana Ringenberg, Dr. John Springer

Social engineering attacks, especially trust exploitation, have become a focus of attention for cybercriminals attempting to manipulate or deceive users to take actions that further expose their vulnerabilities. This has also become a budding field for researchers as these interactions are based on complex social equations that are constantly taken advantage of. Identifying the “weakest link” is a popular method of identifying how these exploits take place, generally by observing when individuals fall for a social engineering attack. However, valuable insights may be used to harden security by observing patterns in users resistant or vigilant to these attacks. Primarily, this trend may be discovered in resistant users’ personality traits. This has been found to be a more accurate indicator of behavior than self-reported intentions. Survey responses (n=120) indicate correlations between high test scores in trust exploitation exercises and Conscientiousness in the Big 5 Personality Model (p<0.001). No significant correlation was seen between self-reported cybersecurity habits and actual security behavior.

NETWORK SECURITY

33. ASMprofiler

Kyle Harvey, Nicholas Bogan, Dr. Gustavo Rodriguez-Rivera

The purpose of this research project was to create a user-friendly profiler that aids students in the optimization of x86_64 Assembly and C. This research project utilizes the profil() system call inside a shared library to create a “histogram” of time spent at each address within the text section of the target program. This data is then read from a file by a Python3 program and combined with data from objdump, addr2line, and the source code of the target program. The source code information allows for compiler-generated Assembly code to be associated with its C source code. This data is then displayed to the user in their browser via a locally hosted anonymous web server. This program successfully determined where the target program was spending its time and displayed it to the user in an easy-to-read format. This program assisted more than 500 CS250 Computer Architecture students in optimizing a hashtable written in Assembly. Computer Science educators may consider utilizing this program to aid students in optimizing C and Assembly code. Programmers may consider utilizing this program to aid themselves in optimizing C with inline Assembly code. This program could also be used for analyzing code at the assembly level to identify the most common execution paths in network exposed programs that open them up to DDoS attacks.

34. Attacking and Improving the Tor Directory Protocol

Zhongtang Luo, Adithya Bhat, Kartik Nayak, Dr. Aniket Kate

The Tor network enhances clients’ privacy by routing traffic through an overlay network of volunteered intermediate relays. Tor employs a distributed protocol among nine hard-coded Directory Authority (DA) servers to securely disseminate information about these relays to produce a new consensus document every hour. With a straightforward voting mechanism to ensure consistency, the protocol is expected to be secure even when a minority of those authorities get compromised. However, the current consensus protocol is flawed: it allows an equivocation attack that enables only a single compromised authority to create a valid consensus document with malicious relays. Importantly the vulnerability is not innocuous: We demonstrate that the compromised authority can effectively trick a targeted client into using the equivocated consensus document in an undetectable manner. Moreover, even if we have archived Tor consensus documents available since its beginning, we cannot be sure that no client was ever tricked. We propose a two-stage solution to deal with this exploit. In the short term, we have developed and deployed TorEq, a monitor to detect such exploits reactively: the Tor clients can refer to the monitor before updating the consensus to ensure no equivocation. To solve the problem proactively, we first define the Tor DA consensus problem as the interactive consistency (IC) problem from the distributed computing literature. We then design DirCast, a novel secure Byzantine Broadcast protocol that requires minimal code change from the current Tor DA code base. Our protocol has near-optimal efficiency that uses optimistically five rounds and at most nine rounds to reach an agreement in the current nine-authority system.

Our solutions are practical: our performance analysis shows that our monitor can detect equivocations without changing the authorities' code in five minutes; the secure IC protocol can generate up to 500 consensus documents per hour in a real-world scenario. We are communicating with the Tor security team to incorporate the solutions into the Tor project.

35. Centralized Hierarchical Cybersecurity Monitoring Towards Securing the Defense Industrial Base Supply Chain

Vijay Sundararajan, Dr. J. E Dietz

With the rise in cyberattacks by nation-state adversaries, the US Department of Defense (DoD) introduced mandatory cybersecurity compliance to fortify the Defense Industrial Base (DIB) supply chain and communication with its private partners. These private partners, obligated by Defense Federal Acquisition Regulations (DFARS), were required to conform to the latest standards in computer and data security. The Cybersecurity Maturity Model Certification (CMMC) is a compliance regulation built upon the existing DFARS 252.204-7012 and the NIST SP 800-171 security controls. These private partners, also referred to as contractors, currently encounter challenges in implementing and monitoring these controls on their information systems, which store, process, and transmit Controlled Unclassified Information (CUI). Safeguarding CUI and confidential communications throughout the supply chain, from the DoD to its contractors/sub-contractors, is imperative to mitigate cyber threats. This paper introduces a centralized hierarchical cybersecurity monitoring (CHCM) model for realtime compliance maintenance. The model is applicable to any type of supply chain relying on information systems to transfer important information and data. Results showcasing the effectiveness of CHCM have been compiled from nine DoD contractors. Furthermore, this paper examines current work on centralized cybersecurity models, elucidates the security aspects, and addresses the benefits and challenges of implementing the CHCM model.

36. E-XAI: Evaluating Black-Box Explainable AI Frameworks for Network Intrusion Detection

Osvaldo Arreche, Tanish Guntur, Jack Roberts

The exponential growth of intrusions on networked systems inspires new research directions on developing artificial intelligence (AI) techniques for intrusion detection systems (IDS). In particular, the need to understand and explain these AI models to security analysts (managing these IDS to safeguard their networks) motivates the usage of explainable AI (XAI) methods in real-world IDS. In this work, we propose an end-to-end framework to evaluate black-box XAI methods for network IDS. We evaluate both global and local scopes for these black-box XAI methods for network intrusion detection. We analyze six different evaluation metrics for two popular black-box XAI techniques, namely SHAP and LIME. These metrics are descriptive accuracy, sparsity, stability, efficiency, robustness, and completeness. They cover main metrics from network security and AI domains. We evaluate our XAI evaluation framework using three popular network intrusion datasets and seven AI methods with different characteristics. We release our codes for the network security community to access it as a baseline XAI framework for network IDS. Our framework shows the limitations and strengths of current black-box XAI methods when applied to network IDS.

37. Explainability of Machine Learning in Intrusion Detection Systems

Yujie Zhang, Xianshun Jiang Advised by Dr. Mohammad Noureddine

With the proliferation of network security attacks, various kinds of intrusion detection systems have been invented and so as to deal with new emerging threats, machine learning techniques are widely used on network intrusion detection systems (NIDS). However, problems still exist for machine learning models in NIDS such as the existence of semantic gaps, shortcut learning, and the high cost of errors. In order to better understand and improve the model, we need to utilize the explainability of machine learning models. This paper provides the explainability of models from both local and global sides based on the CIC-Bell-DNS-EXF-2021 dataset, and uses shapely values, an idea from game theory, to help understand how features affect the prediction result. Besides, this paper also focuses on analyzing the distribution of SHAP values (SHapley Additive exPlanations) of different features and corresponding values in order to generate a formula, which gives weight to features and values, for the firewall rule to better detect the malicious incoming traffic.

38. Global and Distributed Reproduction Numbers of a Multilayer SIR Model with an Infrastructure Network

Jose I. Caiza, Junjie Qin, Dr. Philip E. Pare

In this paper, we propose an SIR virus model in a user network coupled with an infrastructure network that has a virus spreading in it. We develop a threshold condition to characterize the monotonicity and peak time of a weighted average of the infection states in terms of the global (network-wide) effective reproduction number. We further, define the distributed reproduction numbers (DRNs) of each node in the multilayer network which are used to provide local threshold conditions for the dynamical behavior of each entity. Furthermore, we leverage the DRNs to predict the global behavior based on the node-level assumptions. We use both analytical and simulation results to illustrate that the DRNs allow a more accurate analysis of the networked spreading process than the global effective reproduction number.

39. Leo: Online ML-based Anomaly Detection at Multi-Terabit Line Rate

Syed Usman Jafri, Dr. Sanjay Rao, Vishal Shrivastav and Mohit Tawarmalani

Online traffic classification enables critical applications such as network intrusion detection and prevention, providing Quality-of-Service, and real-time IoT analytics. However, with increasing network speeds, it has become extremely challenging to analyze and classify traffic online. In this paper, we present Leo, a system for online traffic classification at multi-terabit line rates. At its core, Leo implements an online machine learning (ML) model for traffic classification, namely the decision tree, in the network switch's data plane. Leo's design is fast (can classify packets at switch's line rate), scalable (can automatically select a resource-efficient design for the class of decision tree models a user wants to support), and runtime programmable (the model can be updated on-the-fly without switch downtime), while achieving high model accuracy. We implement Leo on top of Intel Tofino switches. Our evaluations show that Leo is able to classify traffic at line rate with nominal latency overhead, can scale to model sizes more than twice as large as state-of-the-art data plane ML classification systems, while achieving classification accuracy on-par with an offline traffic classifier.

40. Rationality of Learning Algorithms in Repeated Games

Shivam Bajaj, Pranoy Das, Yevgeniy Vorobeychik, Dr. Vijay Gupta

Many learning algorithms are known to converge to an equilibrium for specific classes of games if the same learning algorithm is adopted by all agents. However, when the agents are self-interested, a natural question is whether agents have a strong incentive to adopt an alternative learning algorithm that yields them greater individual utility. We capture such incentives as an algorithm's rationality ratio, which is the ratio of the highest payoff an agent can obtain by deviating from a learning algorithm to its payoff from following it. We define a learning algorithm to be c-rational if its rationality ratio is at most c irrespective of the game. We first establish that popular learning algorithms such as fictitious play and regret matching are not c-rational for any constant c. We then propose and analyze two algorithms that are provably 1-rational under mild assumptions, and have the same properties as (a generalized version of) fictitious play and regret matching, respectively, if all agents follow them. Finally, we show that if an assumption of perfect monitoring is not satisfied, there are games for which c-rational algorithms do not exist, and illustrate our results with numerical case studies.

41. RiFT: Cyber Adversary Likelihood

Phoebe Abbruzzese, Madhu Joshi, Gabe Samide, Dr. Courtney Falk, Dr. Rick Kennell

The Cyber Adversary Likelihood project has the goal of identifying methods for modeling adversaries in attacks on critical infrastructure and using those models to help determine the likelihood of various adversary actions. Specifically, the project will examine adversary actors in the context of cyber systems (including information systems and control networks) and propose modeling approaches to approximate their behaviors. The project will develop a method to estimate likelihoods of various adversary actions in relevant contexts and then characterize and demonstrate that method. The ultimate use case of the model(s) and tool(s) is to estimate likelihood parameters in a broader model that will be used to assess risk to critical infrastructure from malicious and non-malicious hazards.

42. Sharding SMR with Optimal-size Shards for Highly Scalable Blockchains

Jianting Zhang, Zhongtang Luo, Raghavendra Ramesh, Dr. Aniket Kate

Blockchain relies on State Machine Replication (SMR) to enable trustless nodes to uphold a consistent ledger while tolerating Byzantine faults. With the rapid growth of decentralized web3 platforms and applications, a central challenge of blockchain systems is scalability, which can be evaluated with two metrics: high performance and large network. However, existing blockchain systems struggle to simultaneously achieve both scalability metrics while requiring to guarantee the underlying security properties of SMR—safety and liveness. In this poster, we present a novel blockchain architecture addressing this dilemma by sharding the SMR. Our architecture builds upon two core insights: ordering-processing sharding scheme and safety-liveness separation. Specifically, the ordering-processing sharding scheme securely accommodates a large number of nodes by dividing them into multiple shards, enhancing the network scale. Additionally, the safety-liveness separation allows each shard to consider the security properties of SMR against Byzantine failures separately, by which the system can create more optimal-size shards to process transactions in parallel, enhancing performance. The preliminary experiments show the efficacy of our architecture in scaling blockchains.

43. Snooping Pay-over-the-Phone Transactions over Encrypted 5G/4G Voice Calls

Jingwen Shi, Shaan Shekhar, Guan-Hua Tu, Dr. Chunyi Peng

In this poster, we present a new attack to snoop pay-over-the-phone transactions over encrypted 5G/4G voice calls. We deploy a radio sniffer to eavesdrop 5G/4G communication, followed by inferring confidential pay-over-the-phone transaction despite encryption protection. Interactive Voice Responsive (IVR) technology and other 3GPP standards for enhancing 5G/4G calls widely adopted by mobile network operators makes the attack feasible. While these enhancements enable detection of voice calls over the 5G/4G traffic, IVR-specific features help us identify the presence of an IVR call, followed by spying of sensitive payment transactions over the 5G/4G traffic in real-time.

POLICY, LAW AND MANAGEMENT

44. A Cybersecurity Testbed for Connected and Autonomous Vehicle Systems

Zengxiang Lei, Ruichen Tan, Dr. Satish V. Ukkusuri

The Transportation Cybersecurity and Resilience Center (TraCR) is developing a cybersecurity testbed for connected and autonomous vehicle systems. The goal of this testbed is to comprehensively evaluate the impact of cyberattacks on CAV systems and assess the efficacy of defense algorithms. To achieve this, the testbed incorporates the full life cycles of transportation services and realistic data streams. Specifically, it encompasses data collection, prediction, decision generation, and actuation across various transportation services (including private vehicles, ride-hailing, and public transit), with the data stream modeled using an open-source stream-processing platform named Kafka. This poster shows the testbed design and current development progress. The specific use cases are also highlighted.

45. Adversarial booking attack for autonomous on-demand mobility services

Zengxiang Lei, Dr. Satish V. Ukkusuri

On-demand mobility services, such as Uber and Lyft, are at the forefront of transforming transportation operations by providing online vehicle scheduling and routing. Recently, fully controlled fleets have also been realized through autonomous driving. Despite the clear benefits of introducing real-time controls in transportation operations, the vulnerabilities and associated risks are largely understudied. In this study, we investigate a new attack model named "adversarial booking attacks" to measure the risks inherent in the core operation of on-demand mobility services--the request-vehicle matching process. The attack involves malicious entities who manipulate multiple accounts to generate purposefully crafted trip requests, aiming to disrupt service operations. We formulate the attack into an optimization problem with objectives to reduce the matching pairs and maximize induced traffic in a specific region. Using real-world ride-hailing trip records from New York City and traffic simulator SUMO, we explore the potential large-scale impact of adversarial booking attacks under various demand scenarios and attack strategies.

46. Securing the Future: A Strategic Framework for Cyber Liability Insurance

Janith D'Alwis

Abstract: With cybercrimes predicted to cost the world $10.5 trillion in 2025, this cyber liability framework is a crucial blueprint for businesses navigating the ever changing digital landscape. As cybercrime continues to become more sophisticated the financial repercussions for unprepared entities surge, underscoring the indispensable role of cyber liability insurance. This research delineates a comprehensive insurance framework tailored to mitigate financial losses, safeguard digital assets, and ensure business continuity. By incorporating a meticulous risk assessment, customized coverage options, and promoting cybersecurity best practices, the proposed framework addresses the multifaceted nature of cyber risks. This research provides a strategic pathway for businesses to shield themselves against the burgeoning financial and operational impacts of cyber threats, thereby securing their future in an increasingly digitized global landscape.

PREVENTION, DETECTION AND RESPONSE

47. Code Blue - Gamification of Incident Response

Nicole Hands, Olivia Anderson

Gamification, or the process of applying game-like elements to non-game scenarios, has been shown to increase learners' skill and knowledge mastery, motivation, interest, and enjoyment. Applying gamification to cybersecurity incident response would allow learners to develop the technical and decision-making skills required to effectively detect and respond to cyber incidents in a more accessible and appealing format. This project looks at improving cybersecurity education through gamification. The solution we've developed utilizes Roguelike games to provide a randomized scenario in which students can practice incident response in real-world like conditions.

48. Forensics Analysis of Oura Ring Gen 3 on Android, iOS and Cloud

Xiao Hu, Akif Ozer, Miloš Stanković and Dr. Umit Karabiyik

The popularity, variety, and accessibility of wearable devices have surged, with sales exceeding 490 million units in 2022 alone. This trend offers advantages to consumers but presents significant challenges for mobile forensic experts. These devices collect and manage extensive user data, from health metrics to personal details, creating a potential treasure trove for legal evidence. Consequently, forensic investigators must continually adapt and expand their expertise to navigate this evolving landscape. Among these technologies, smart rings, particularly the Oura Ring Gen 3, have attracted significant attention within the forensic community. This paper examines the Oura Ring Gen 3 application, or Oura, through a forensic lens, analyzing data across Android, iOS, and Cloud platforms. It details the specific data paths and locations, offering insights that could be crucial for digital forensic investigations and informing users about the nature of the information stored by these devices. Essential data points such as heart rate, activity types, user ID, and other personal details are highlighted for their potential value in various contexts.

49. Light Curve Shape Inversion

Liam Robinson, Dr. Carolin Frueh

Many areas of Space Domain Awareness (SDA) require shape information, but such information is not readily available. With an optical telescope, we can collect resident space object (RSO) brightness data over time, known as a light curve. Since the light curve is a product of shape, attitude, and materials of the object, how can we recover those attributes?

50. MIXED-SENSE: A Mixed Reality Sensor Emulation Framework for Test and Evaluation of UAVs Against False Data Injection Attacks

Kartik Anand Pant, Li-Yu Lin, Jaehyeok Kim, Worawis Sribunma, Dr. James Goppert

We present a high-fidelity Mixed Reality sensor emulation framework for testing and evaluating the resilience of Unmanned Aerial Vehicles (UAVs) against false data injection (FDI) attacks. The proposed approach can be utilized to assess the impact of FDI attacks, benchmark attack detector performance, and validate the effectiveness of mitigation/reconfiguration strategies in single-UAV and UAV swarm operations. Our Mixed Reality framework leverages high-fidelity simulations of Gazebo and a Motion Capture system to emulate proprioceptive (e.g., GNSS) and exteroceptive (e.g., camera) sensor measurements in real-time. We also propose an empirical approach to faithfully recreate signal characteristics such as latency and noise in these measurements.

51. Risk Assessment of Multi-Agent System Under Denial-of-Service Cyberattacks Using Reachable Set Synthesis

Minhyun Cho, Sounghwan Hwang

Multi-agent systems (MASs) have vulnerabilities to various types of cyberattacks disrupting inter-agent communication. To assess the potential risk associated with these cyberattacks, this paper proposes a proactive risk assessment method using reachable set synthesis. Denial-of-Service (DoS) attacks, where adversaries can disrupt communication by a sequence of link disconnections (dynamic alterations), are specifically considered. Our method employs the calculation of reachable sets using Lyapunov functions and linear matrix inequalities (LMIs) derived from them. The proposed method can evaluate the risk of DoS attacks for both individual agents and the entire system in two levels by computing over-approximated ellipsoidal reachable sets. To demonstrate the applicability of our method, we provide an illustrative example involving a leader-follower MAS performing formation control in an adversarial environment with scattered obstacles.

52. Safety-Critical Control for Nonlinear Affine Systems with Robustness and Attack Recovery

Sungsoo Kim, Minhyun Cho, Sounghwan Hwang

We propose a safety-critical controller design for nonlinear affine systems under actuator cyberattacks and model uncertainties. To achieve this, our approach addresses two primary challenges: First, we propose a robust controller that employs a sliding mode-based control barrier function (SM-CBF) to adeptly manage model uncertainties. Second, we devise a CBF-based attack detection mechanism to promptly alert the presence of cyberattacks within the actuator channel. The natural drawback of a conventional CBF-based controller is that its performance (safety guarantees) significantly depends on the model uncertainties. The inclusion of these uncertainties may lead to safety breaches/violations. To overcome this technical challenge, our proposed controller with SM-CBF approach enables us to effectively enforce the safety conditions of the system even in the presence of model uncertainties. Furthermore, we propose a novel CBF-based attack detector such that it can determine whether the system trajectory moves toward the outside of the safety set by cyberattacks. Through the synthesis of the designed robust controller and attack detector, our proposed safety-critical controller can greatly enhance system safety while concurrently achieving control performance. Finally, an illustrative example of the stabilization of quadrotor unmanned aerial vehicles (UAVs) is provided to demonstrate the effectiveness of the proposed methodology.

53. SiDG-ATRID: Simulator for Data Generation for Automatic Target Recognition, Identification and Detection

Younggil Chang, Alec Andrulis, Isabel Hoppe

The increased utilization of Unmanned Aerial Vehicles (UAVs) in diverse missions, from humanitarian aid to combat operations, underscores the necessity for an efficient and cost-effective development workflow for autonomous systems. Especially for defense purposes, building autonomous target recognition systems capable of detecting, identifying, and classifying adversarial agents with machine learning models requires extensive data for training. Consequently, simulation software has become an essential tool for developers seeking to assess autonomous system performance and collect data across various environments. Furthermore, the transition to real-world, application-ready systems necessitates a simulation platform that replicates not only the vehicle control algorithms but also environmental factors that affect system performance, such as lighting conditions and sensor noise. In response to these requirements, we introduce 'SiDG-ATRID' (Simulator for Data Generation for Automatic Target Recognition, Identification and Detection), a simulation platform that enables the collection of high-fidelity imagery data, powered by Unreal Engine 5. The simulator supports multi-agent simulations using the AirSim API library for UAV controls and simulates commercial aircraft traffic. This framework allows for customized camera placements to record videos or photos and manage environmental conditions such as weather and lighting. Additionally, by leveraging the Cesium API for geospatial mapping, it can accurately recreate real-world environments, enhancing the realism and applicability of simulations. This integrated approach enhances the efficiency and effectiveness of synthetic data generation for detection tasks, enabling developers to easily configure simulations and collect diverse data.

54. Technology and the Emerging Force of Change: Cyber Secure Competency Framework for Older Adults Using Delphi Method

Julie Wenner

The United Nations (UN) and World Health Organization (WHO) have teamed up by identifying the decade of 2021- 2030 as “Healthy Ageing, which requests the whole of government and society response to seek to do business differently towards transformative and positive changes for senior adults, their families, and communities” (WHO, n.d.). Stating that “longer lives are one of humanity’s greatest achievements and ensuring there is equal access to technologies across all ages” (WHO, n.d.; Orłowska & Błeszyńska, 2020, Abstract). This study will address the problem that digital technology users incur cyber assaults from unknown and known nefarious actors (Verizon, 2022). To date, Nation States have not included the element of training their citizens on how not to become a statistic within the digital cyberspace of the internet—leaving the general populace uninformed, lacking secure cyber competencies, framework, and knowledge not to become a subsequent digital exploitation. The literature review identified that Cyber Secure Core Competencies could be designed to create a human firewall methodology mindset for senior adults to change their mental model of Digital Information Systems Protections (ISP). This research proposal introduces security awareness to improve senior adults’ poor digital hygiene to create a change in behavior with perceived security awareness; this proposed research study will focus on the following three questions: Question 1 [RQ1]: What competency skills and technical knowledge do senior adult digital users need to adopt to protect themselves online to minimize potential cyber-attacks based on common cybersecurity threats and risks? Question 2 [RQ2]: What are the attitudes, barriers, and challenges that senior adults encounter in adopting cybersecurity best practices? Question 3 [RQ3]: What Cybersecurity practices from other fields, such as digital accessibility or technology adoption and existing Digital Competency models, can be adapted to create a Cybersecurity Competency Model for Senior Adults? The proposed research will begin in May 2023, with a proposed completion in December 2024. Additional literature reviews of core competency models, the Delphi method design criteria structure for the expert panel member selection, and the Purdue University, West Lafayette, IN, IRB requirements are necessary for this research proposal. The best defense against Cyber-attacks is a practiced offense.

55. The Power of Digital Forensics in Smart Device Investigations

Miloš Stanković, Xiao Hu, Dr. Umit Karabiyik, Dr. Marcus K. Rogers, Dr. Smriti Bhatt

As an important branch of forensic science, digital forensics is growing in importance, with mobile forensics being a top priority, focusing on the collection of data from mobile devices such as smartphones, tablets, and GPS systems. These ubiquitous mobile devices are treasure troves of stored user information, covering call logs, text message exchanges, use of social platforms, location data, and so on. While becoming an important part of people's lives, these devices are also becoming vehicles for criminal activity. However, as technology continues to grow, forensic studies on these devices have seen significant gaps. In response to these related challenges, this study created two real-world scenarios that can be applied to real-world digital investigations. Based on these real-world crime case scenarios, multiple smart devices were utilized to determine the ins and outs of the necessary digital evidence associated with each event, thereby demonstrating the investigative process from start to finish. Our findings included the analysis of identifiable and relevant data and related artifacts designed to aid in digital forensic investigations from multiple PIoT devices. Additionally, we pointed out the differences in storing data between devices and the cloud for digital forensic investigations. Finally, we created a publicly available database containing the findings and vulnerabilities found in the devices we have used.

56. Using Digital Twins as a Sandbox for the Evaluation of Cyber Attacks on Avionics Networks

Alisha Gadaginmath, Sanjana Gadaginmath, Yury A. Kuleshov, Hridhay Monangi (TA), Kabir Nagpal, Katie B. O'Daniel, Dalbert Sun, Lucas Tan, Korel Ucpinar, Nathan L. Veatch, Naren Velnambi , Dr. Mark Daniel Ward

Conventional methods of the evaluation of cyber attacks on avionics networks do not meet the requirements of the ongoing Industry 4.0 (I4.0) revolution. The proposed alternative method is the use of digital twins. Digital twins have not been widely used in avionics networks in academia. During the previous stages of research, the Purdue Data Mine students in collaboration with the Boeing Company mentors developed a prototype of a digital twin and simulated sample attacks. Based on the recent feedback from aviation industry experts, the research team is currently focused on increasing the fidelity of the digital twin and integrating the new attack vector related to Automatic Dependent Surveillance-Broadcast (ADS-B). The project goal is to design a digital twin-based avionics networks sandbox and validate its features with the new attack vector. Other students and industry researchers can potentially benefit from using the sandbox to evaluate cyber attacks on avionics networks.

57. Zero Trust Chain (ZTC): Security Solutions for 5G Networks with an O-RAN-Centric and Device-Centric Approach

Yongkyu Jang. Dr. David Love, Dr. Taejoon Kim, Dr. Christopher Brinton, Dr. Sonia Fahmy, Dr. Remi Chou, Dr. Vuk Marojevic, Dr. Syed Rafiul Hussain, Dr. Hyuck Kwon, Dr. Sang W. Kim

Our team, named Zero Trust X (ZTX), proposes a software solution that enables military squads to securely share situational awareness in their operations through high-performance, but often untrusted, 5G networks. It will allow DoD operators to discover malicious entities in near-real-time and provide communication mechanisms to avoid adversary's control over DoD traffic. Specifically, through a minimum amount of cooperation with the network operator, part of our solution leverages O-RAN for new threat monitoring and mitigation solutions specifically designed for 5G networks. We complement this O-RAN-centric approach with a device-centric approach to ensure that DoD devices also implement their own layer of security and do not solely rely on the security protocols of the network provider that could possibly be compromised by adversaries. Such a combination will substantially enhance the security of the whole system. Importantly, our device-centric solutions do not require cooperation from the network providers, nor any changes to the existing 5G standards.

About CERIAS

CERIAS — The Center for Education and Research in Information Assurance and Security — is the world's largest and foremost multidisciplinary academic institute addressing the issues cyber and cyber-physical security, assurance, privacy, forensics, artificial intelligence, and trusted electronics. CERIAS brings together a team of world-class faculty, graduate student researchers and industry partners with the shared goal of advancing the state of cyber security through basic and applied research. CERIAS serves as an unbiased resource of information to the worldwide community.

Faculty from eight different colleges, and more than 18 departments, across Purdue University are conducting CERIAS research. The six primary areas of CERIAS research are:

- Assured Identity and Privacy
- End System Security
- Human Centric Security
- Network Security
- Policy, Law and Management
- Prevention, Detection and Response

Research at CERIAS continues to be vibrant with current projects addressing a large number of topics, from networks, operating systems, and database security to forensics and human factors. Security research at CERIAS results in comprehensive approaches and is characterized by both theoretical and experimental results. Notable efforts are also devoted to the development of testbeds and experimental environments; examples include the SOL4CE Laboratory, VoIP testbed, the Biometrics Laboratory and the ReAssure system. Education of top security researchers is a key goal of CERIAS - and students (undergraduate, graduate and post-doctoral) are involved in all those projects. We trust that you will appreciate this sampler of our projects.

Detailed information about research being conducted at CERIAS or at one of our academic partners is available by contacting us at (765) 494-7841ß or by visiting www.cerias.purdue.edu.

CERIAS has Moved to the Discovery Park District!

The *CERIAS Galactic Headquarters* has moved across campus to the Convergence Center for Innovation and Collaboration (CONV). Come visit us!

CERIAS, Purdue University
101 Foundry Drive
Convergence Center
Suite 3800
West Lafayette IN 47906-3446

Just south of Mitch Daniels Boulevard
(formerly State St.) on the west side of campus.

LOCAL RESTAURANTS

Provided by Purdue Conferences

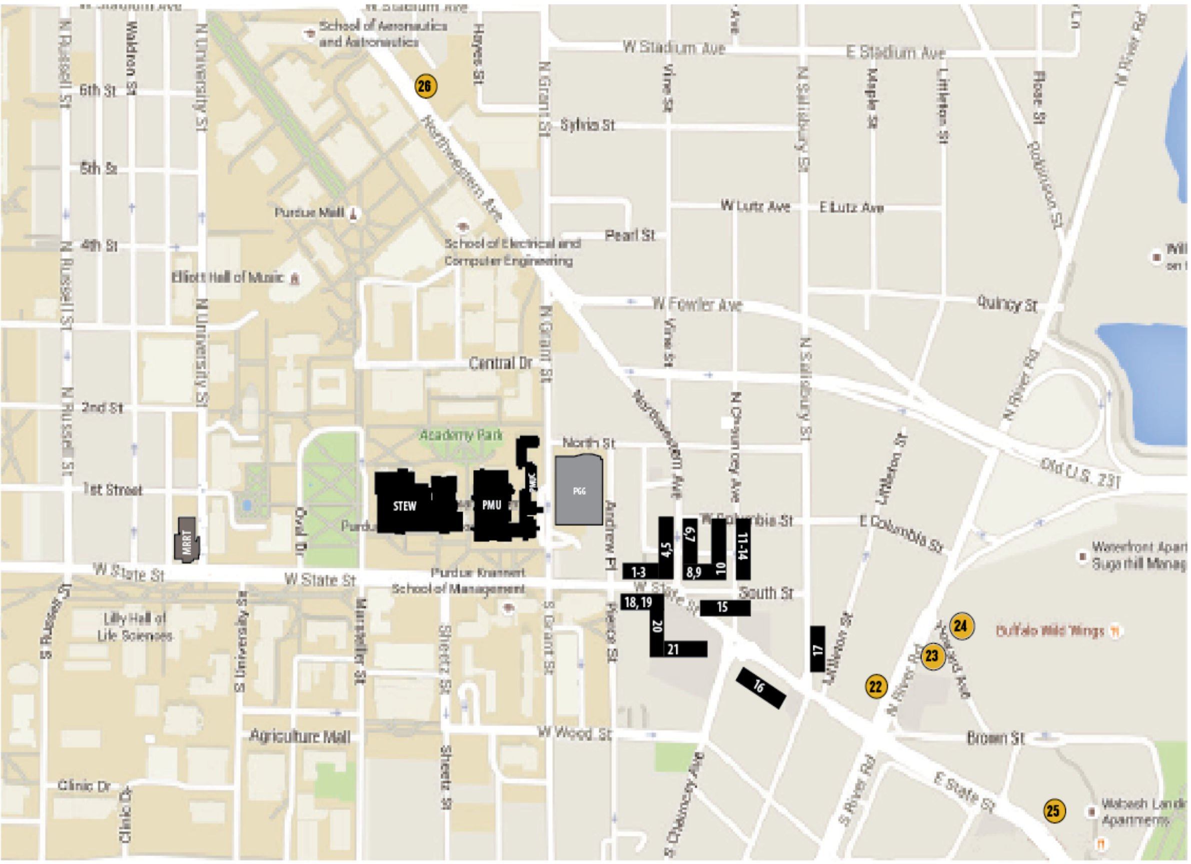

ON CAMPUS

PURDUE MEMORIAL UNION (PMU)

LOWER LEVEL

Purdue Memorial Union's newly renovated space has a wide variety of dining options. From Starbucks, to Sushi, to Burgers, Pizza, Mexican, and more...visit: www.union.purdue.edu/dine/

SECOND FLOOR

8 Eleven Modern Bistro

STEWART CENTER (STEW)

Newsstand

MARRIOT HALL (MRRT)

Boiler Bistro
Lavazza

NEARBY

1. Mad Mushroom
2. Brothers
3. Blue Nile
4. Potbelly Sandwiches
5. Einstein Bros. Bagels
6. Panda Express
7. Egyptian Café
8. Greyhouse Coffee
9. Vienna Expresso Bar
10. Majé Sushi
11. Maru Sushi
12. Fiesta Mexican Grill
13. Red Mango
14. Noodles & Company
15. Chipotle
16. Raising Cane's Chicken Fingers
17. Triple XXX
18. Harry's
19. Jimmy Johns
20. Five Guys Burgers
21. Basil Thai & Bubble Tea
22. Town & Gown Bistro
23. Nine Irish Brothers
24. La Hacienda Bar & Grill
25. Moe's
26. Another Broken Egg

Stewart Center Wireless Information

For Purdue Students, Staff and Faculty:

- Use any of the following SSIDs: ‘PAL3.0’ or ‘eduroam’.
- Login with your Purdue career account credentials.

For Visitors:

- Connect to the ‘attwifi’ SSID
- Open your web browser (Firefox, Chrome, IE, etc.)
- Click on the **“Get Connected”** button.

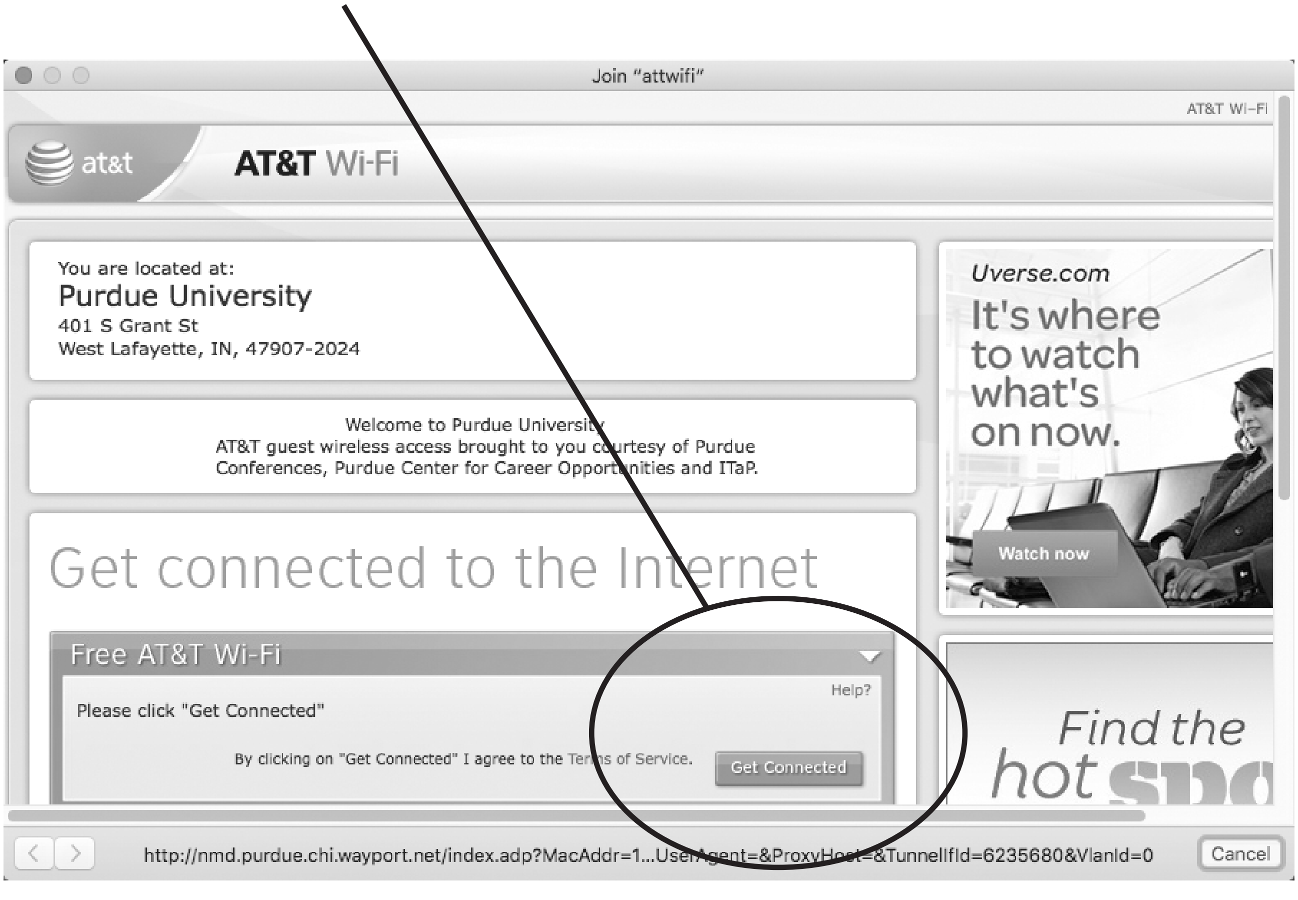

You are located at:
Purdue University
401 S Grant St
West Lafayette, IN, 47907-2024

Welcome to Purdue University
AT&T guest wireless access brought to you courtesy of Purdue Conferences, Purdue Center for Career Opportunities and ITaP.

Get connected to the Internet

Free AT&T Wi-Fi

Please click "Get Connected"

By clicking on "Get Connected" I agree to the Terms of Service.

http://nmd.purdue.chi.wayport.net/index.adp?MacAddr=1...UserAgent=&ProxyHost=&TunnelIfId=6235680&VlanId=0