# FALL 2025 RESEARCH REPORTS


PURDUE UNIVERSITY®
Center for Education and Research in Information Assurance and Security (CERIAS)

www.cerias.purdue.edu

# Table of Contents

**Introduction** ..... 12

**What's New (Fall 2025)** ..... 13

**Recently Added (Spring 2025)** ..... 14

**Artificial Intelligence and Machine Learning** ..... 15

Average Reward Reinforcement Learning ..... 15

Intelligent Transportation Systems ..... 16

Machine Learning Approach to Joint Multi-Agent Decision Making ..... 17

Optimal Sample Complexity Guarantees for Discounted Reward Reinforcement Learning ..... 18

Quantum Machine Learning ..... 18

Reinforcement Learning for Network Resilience ..... 19

Submodular Bandits ..... 20

Temporal Abstractions in Multi-Agent Reinforcement Learning ..... 21

Malware Speaks! Deep Learning Based Assembly Code Processing for Detecting Evasive Cryptojacking ..... 22

DARPA Project GNOME with USC ..... 23

Security and Privacy of Large Language Models (LLMs) ..... 23

FAI: Identifying, measuring, and mitigating fairness issues in AI ..... 24

Exploring the Process, Challenges, and Effective Practices of Reproducing and Adapting Machine Learning Models ..... 24

Reusing Deep Learning Models: Challenges and Directions in Trustworthy Software Engineering ..... 25

Using GenAI in Software Engineering Education: Efficiency and Cybersecurity ..... 27

Pose and Range Estimation from Monocular Monstatic Single Resolved Images ..... 27

Automatic Discovery of Real-Time Visual Data on the Internet ..... 28

NSF CRII: RI: A Study of Rank-based Decomposable Losses for Machine Learning ..... 28

Active Learning using Loss Geometry ..... 29

Machine Unlearning ..... 29

Scalable Data Efficient Learning ..... 29

Artificial Intelligence for Music ..... 29

Computer Vision using Contextual Information ..... 30

Dynamic Optimization Using Object Detection in Bandwidth Constrained Automatic Multi-camera Networks ..... 31

Low-Power Computer Vision ..... 31

Securing Data Privacy at the Edge Using Trusted Execution Environment ..... 32

Development of a Reliable Method for General Aviation Flight Phase Identification ..... 32

AI/Control for Intrusion Detection in Constrained Embedded Systems ..... 33

Phishing Email Detection Through Machine Learning and Word Error Correction ..... 33

Data Preprocessing and ML Model Fairness ..... 34

Explainable AI: Foundations, Applications, Opportunities for Data Management Research ..... 34

Generating Interpretable Data-Based Explanations for Fairness Debugging ..... 35

Generating Interpretable Diagnostic Explanations for Black-box AI Systems ..... 36

Artificial Infrastructures ..... 37

Testing Deep Learning Software ..... 38

RNCP: A Resilient Networking and Computing Paradigm for NASA Space Exploration ..... 38

Generating Electricity Managed by Intelligent Nuclear Assets (GEMINA) ..... 39

EAGER: SaTC-EDU: Cybersecurity Education in the Age of Artificial Intelligence: A Novel Proactive and Collaborative Learning Paradigm ..... 40

Probabilistic Red-Teaming for Large Language and Vision-Language Models ..... 41

Inter-Play Between Cyberscurity and Deep-Learning ..... 42

**Assured Identity and Privacy** ..... 43

Directed Infusion of Data ..... 43

Obfuscation of Audio Signals ..... 43

Foundations of Cyber-Physical Infrastructure for Creative Design and Making of Cyber-physical Products ..... 44

Customized Privacy Mechanisms for Statistical Inference ..... 45

Differential Privacy Methods for Machine Learning and Complex Data Structures ..... 46

Foundations of Differential Privacy ..... 46

Statistical Methods for Privatized Data ..... 47

ABAUS: Active Bundle AUthentication Solution Based on SDN for Vehicular Networks ..... 48

An Attack to One-Tap Authentication Services in Cellular Networks ..... 49

Privacy – Preserving Data Dissemination in Untrusted Cloud ..... 50

QPCASIN: A Quantum-Defended Privacy-Aware Preemptive Handover-Enabled Continuous Authentication in Space Information Networks ..... 50

Formal Privacy for Multi-stage Probability Samples ..... 51

Securing the Software Supply Chain: Theories, Measurements, Runtime Defenses, and Software Signing Infrastructure (PKI) for commercial and open-source software ..... 51

Language Support for Precise Privacy-Preserving Computation ..... 53

Standoff Inverse Analysis and Manipulation of Electronic Systems ..... 54

Effects of Body Position on Facial Recognition in Police Body-Worn Cameras ..... 55
Evaluation of Face Recognition for Law Enforcement: Book-In Photo Acquisition and Surveillance Camera Placement ..... 55
Post-Mortem Biometrics ..... 56
A Comprehensive Approach for Data Quality and Provenance in Sensor Networks ..... 56
Differentially Private Data Synthesis: Practical Algorithms and Statistical Foundations ..... 57
Privacy-Preserving Data Publishing ..... 58
CRII: SaTC: Securing Smart Devices with AI-Powered mmWave Radar in New-Generation Wireless Networks ..... 58
Towards Machine-learnable Enhancing Framework for Local Differential Privacy ..... 59
In-Toto: Securing the Software Supply Chain ..... 60
SigStore: a Transparent Software Supply Chain Storage System ..... 61
Towards Formal, Risk Aware Authorization ..... 61
A Novel Approach to Robust, Secured, and Cancellable Biometrics ..... 62
Evaluation of Clinical and Genomic Information Privacy Risks From Inference Attacks ..... 63
Secure Video Stream Framework for Dynamic and Anonymous Subscriber Groups ..... 64
Trusted Medical Information System and Health Informatics ..... 64

**Autonomous Systems ..... 66**

Adversarial Examples against Distributed Machine Learning Algorithms ..... 66
Addressing Safety and Security Challenges in ML-based AV Software Stack – Remote Operation Support and Balancing Trade-offs ..... 67
Collaborative Research: CPS: Medium: Transforming Connected and Automated Transportation with Smart Networking, Cooperative Sensing, and Edge Computing ..... 67
Object Recognition using Light Curve Inversion ..... 68

Observability for Autonomy and Sensor Network Design ..... 68
Satellite Imaging using Compressed Sensing ..... 69
Hybrid UAM Model Checking ..... 69
Further Refinement and Integrated Platform for INDOT Traffic Management and Safety Toolset ..... 70
Development of Verification and Validation System for Security Assessment of ECU Components ..... 70
Real Time CFD Mapping of Hazardous Airflow Around Bridge Infrastructure ..... 71
AI-assisted Dynamic Adaptive Planning for Human-in-the-Loop Multi-Agent Systems ..... 71
Secure and Safe Assured Autonomy ..... 72
Integrating Large-Scale Machine Learning and Edge Computing for Collaborative Autonomous Vehicles ..... 72
Secure and Safe Assured Autonomy ..... 73
Whitebox Testing, Debugging, and Repairing for Multi-module Autonomous Vehicles in Near-Collision Traffic Scenariosv ..... 73

**Cryptology and Rights Management ..... 75**

Self-Healing Images ..... 75
The Garbled Computer: Towards Computing without Seeing ..... 75
Password Hashing Algorithms ..... 76
Decentralized Anonymous Credentials from Blockchains ..... 77
Flexible Anonymous Credentials from zk-SNARKs and Existing Identity Infrastructure ..... 77
Privacy Preserving Software Bill of Materials ..... 78
Removing the Human Element: Securing Deployed Cryptographic Systems through the use of Cryptographic Automation ..... 79
zk-creds: Flexible Anonymous Credentials from zkSNARKs and Existing Identity Infrastructure ..... 80

SaTC: Frontiers: Center for Distributed Confidential Computing (CDCC) ..... 81

**Cyber-Physical Systems** ..... 82

Covert Cognizance ..... 82

Securing Embedded Devices by Enforcing Lowest Privilege Execution ..... 83

Characterization of soft dielectric breakdown in GaN MOSHEMTs ..... 84

Compositional IoT Safety and Security in Physical Spaces ..... 85

Bringing Fuzzing to the Cyber-Physical World ..... 85

Security of Autonomous Vehicles ..... 86

Formal methods and Fuzzing for Security in Internet of Things, Embedded Systems,Real-time Operating Systems, and General Software ..... 86

Promoting Inter- and Intra-Organizational Learning from Software Failures: Towards a Failure-Aware Software Development Lifecycle ..... 88

CAREER: Securing Next-Generation Transportation Infrastructure: A Traffic Engineering Perspective ..... 89

Collaborative Research: SaTC: Medium: Exposing and Mitigating Security/Safety Concerns of CAVs: A Holistic and Realistic Security Testing Platform for Emerging CAVs ..... 90

Cloud-connected Cyber-physical Edge Devices for Electric Vehicles Eco-systems with Low Power and Cost ..... 90

Bootloader Security ..... 91

Securing Deeply Embedded Devices ..... 91

Agricultural Robotic Systems for Greenhouses ..... 92

ONR-BAA: Reactor Simulation Tool for Investigating the Resilience of a Cyberphysical Security Ecosystem ..... 92

**End System Security** ..... 94

Protecting and Securing Supply Chain Data throughout its Lifecycle ..... 94

NARCISSUS: Deriving Correct-By-Construction Decoders and Encoders from Binary Formats ..... 94

MicroVM: Micro Virtual Machines for Managed Languages – Abstraction, contained ..... 95

PeX: A Permission Check Analysis Framework for Linux Kernel ..... 96

Rust for Embedded Systems ..... 96

Assurable Configuration of Security Policies in Enterprise Networks ..... 97

Migrating Enterprises to Hybrid Cloud Architectures ..... 97

Cryptanalysis of RSA ..... 98

Search for Aurifeuillian Factorizations ..... 98

Knowledge Graph Construction for Resilient, Trustworthy, and Secure Software Supply Chains ..... 99

Convicting Exploitable Software Vulnerabilities: Practical Input Provenance-Based Approach ..... 100

Secure Group Communication Over Wired/Wireless Networks ..... 101

**Human Centric Security** ..... 102

Securing IoT-based Cyber-Physical Human Systems against Collaborative Attacks ..... 102

RUDOLF: An Efficient and Adaptive Defense Approach Against Website Fingerprinting Attacks Based on Soft Actor-Critic Algorithm ..... 102

Economics of Password Cracking ..... 103

Security and Privacy in Augmented and Virtual Reality (AR/VR) ..... 104

A Human Factors Perspective on Better Phishing Defenses ..... 104

High-fidelity and Trustworthy Teleinteraction Platform ..... 105

Understanding the Impacts of Human Decision-Making on Security and Robustness of Large-Scale Systems ..... 105

Secure, Composable, & Scalable Framework for Trusted Collaborative Computing ..... 106

Internet Based Electronic Voting Enabling Open and Fair Elections ..... 107

Revocable, Interoperable and User-Centric (Active) Authentication Across Cyberspace ..... 108

**Network Security** ..... 109

Explainable AI Methods for Enhancing AI-Based Network Intrusion Detection Systems ..... 109

A Framework of High-Speed Network Protocol Fuzzing Based on Shared Memory ..... 110

Adaptable Safety and Security in V2X Systems ..... 110

BioKA-ASVN: Biometric-Based Key Agreement Scheme for Air Smart Vehicular Networks Using Blockchain Service ..... 111

Quantitative Evaluation of Extensive Vulnerability Set Using Cost Benefit Analysis ..... 112

Developing a Smart and Fast Semantic Intrusion Detection System ..... 112

CICI: CE: Enhancing Cybersecurity for Broadening Data-Driven Research and Partnerships ..... 112

Router Models and Downscaling Tools for Scalable Security Experiments ..... 113

Building Sophisticated Services with Programmable Anonymity Networks ..... 114

SecureCDN: Providing End-to-End Security in Content Delivery Networks ..... 116

Safeguarding Next-Generation Emergency Services (NG-9-1-1) over Cellular Networks ..... 116

Formalizing Enterprise Firewall Management with Informal and Elastic Specifications ..... 117

Scalable and Resilient Distributed Algorithms for Coordination in Large-Scale Networks ..... 117

Enabling Detection of Elusive Malware by Going Out of the Box with Semantically Reconstructed View (OBSERV) ..... 119

Virtualization-Enabled Malware Research ..... 119

Big Data Security Analyses ..... 120

Development of a Secure and Privacy-Preserving Workflow Architecture for Dynamic Data Sharing in Scientific Infrastructures ..... 121

**Other Security Research** ..... 122

Improving the Security and Usability of the Wear OS Permission Model ..... 122

Investigating and Understanding Digital Bill of Materials ..... 122

Removing the Human Element: Securing Deployed Cryptographic Systems through the use of Cryptographic Automation ..... 123

SGX.Fail: How Secrets Get eXtracted ..... 124

SNARKProbe: An Automated Security Analysis Framework for zkSNARK Implementations ..... 125

**Policy, Law, and Management** ..... 126

Big Data Ethics: detecting bias in data collection, algorithmic discrimination and "informed refusal" ..... 126

Purdue University's Computer and Information Technology program creates framework to support implementing data governance in small and medium enterprises. ..... 127

Assessing Security for Organizations Dealing with At-Risk Populations ..... 128

CICI: RDP: Supporting Controlled Unclassified Information with a Campus Awareness and Risk Management Framework ..... 128

**Prevention, Detection and Response** ..... 129

Subtle Adversarial Intrusion Detection with SONAR Software ..... 129

Behavioral and Game-Theoretic Security Investments in Interdependent Systems ..... 130

An MTD-based Self-Adaptive Resilience Approach for Cloud Systems ..... 131

Autonomous Aggregate Data Analytics in Untrusted Cloud ..... 132

ConFoc: Content-Focus Protection Against Trojan Attacks on Neural Networks ..... 132

Hunting for Insider Threats Using LSTM-based Anomaly Detection ..... 133

Incremental Learning Through Graceful Degradations in Autonomous Systems ..... 133

Machine Learning Models to Enhance the Science of Cognitive Autonomy ..... 134

Quantitative Evaluation of Extensive Vulnerability Set Using Cost Benefit Analysis ..... 134

Scalable Learning Through Error-correcting Codes based Clustering in Autonomous Systems ..... 135

Developing New Mechanisms to Enable Open Source Software (OSS) Supply Chain Transparency ..... 135
Developing Software Sensors for Digital Twin based Cybersecurity ..... 136
System Events and Network Traffic Generation for Realistic Cyber Experimentation ..... 136
A First Look at Third-Party Cyber Threat Hunting ..... 137
Eliminating Regex-based Denial of Service ..... 138
Secured and Safe Assured Autonomy (S2A2) for Urban Air Mobility (UAM) ..... 139
Cybersecurity for Unmanned Aerial System Operations in Urban Environment ..... 141
Scalable and Concurrent Targeted Search for Digital Forensics ..... 142
LLM Assisted Vulnerability Detection ..... 143
Automation of Runway Status Light System ..... 143
Complex Networks and Systems Resilience Against Disruption Propagation ..... 144
Cyber-Collaborative Conflicts and Errors Prevention and Detection for Network Resilience ..... 145
Assessing the Relationship between Hacking and Various Personality Traits ..... 147
Psycholinguistic Automated Detection Tool for Criminal Insiders ..... 147
Cyber Adversary Likelihood Project ..... 148
Use of Deception and Misdirection in Cyber Defense ..... 148
Algorithmic and Graph-Theoretic Approaches to Optimal Sensor Placement in Complex Dynamical Systems ..... 150
Algorithms for Persistent Intelligence, Surveillance, and Reconnaissance by Mobile Platforms ..... 151
Better Static Application Security Testing ..... 151
SafeBet: Secure, Simple, and Fast Speculative Execution ..... 152
Process Coloring: Information-Flow Preserving Approach to Malware Investigation ..... 153

Building an Intelligent, Uncertainty-Resilient Detection and Tracking Sensor Network....154

Precise Calling Context Encoding....154

Causality-Driven Mitigation of Cascading Failures in Distributed Systems....155

Testing and detecting software upgrade failures in data-intensive distributed systems....156

**Security Awareness, Education, and Training....157**

Deploying Cyber Emulation, Modeling, and Analysis Tools on the SOL4CE....157

Teaching and Assessing Threat Modeling Competence in Software Courses using Systems Thinking....157

Crime Scene Surveying for IoT Investigations: National Training and Technical Assistance Program....158

CHEESE: Cyber Human Ecosystem of Engaged Security Education....159

SaTC-EDU: EAGER Enhancing Cybersecurity Education Through a Representational Fluency Model....160

Building an Electronic Voting Technology Inspired Interactive Teaching and Learning Framework for Cybersecurity Education....161

**CERIAS-Affiliated Purdue Laboratories and Center....176**

**CERIAS Strategic Partnership Program....177**

# Introduction

The Center for Education and Research in Information Assurance and Security (CERIAS), a cross-cutting institute at Purdue University, is the world's foremost interdisciplinary academic center for cyber and cyber-physical systems; more than a hundred researchers addressing issues of security, privacy, resiliency, trusted electronics, autonomy and trustworthy artificial intelligence. CERIAS brings together world-class faculty, students and industry partners to design, build and maintain trusted cyber/cyber-physical systems. CERIAS serves as an unbiased resource to the worldwide community.

The 11 primary areas of CERIAS research are:

- Artificial Intelligence and Machine Learning
- Assured Identity and Privacy
- Autonomous Systems
- Cryptology and Rights Management
- Cyber-Physical Systems
- End System Security
- Human Centric Security
- Network Security
- Policy, Law and Management
- Prevention, Detection and Response
- Security Awareness, Education, and Training

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 test beds and experimental environments; examples include the SOL4CE Lab, VoIP test bed, 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 for any of the enclosed summarized projects, as well as additional research being conducted at CERIAS or at one of our academic partners, is available by contacting us at (765) 494-7806 or by visiting www.cerias.purdue.edu.

# What's New (Fall 2025)

### Artificial Intelligence and Machine Learning

- Security and Privacy of Large Language Models (LLMs) (pg. 23)
- Active Learning using Loss Geometry (pg. 29)
- Machine Unlearning (pg. 29)
- Scalable Data Efficient Learning (pg. 29)
- Probabilistic Red-Teaming for Large Language and Vision-Language Models (pg. 41)

### Assured Identity and Privacy

- An Attack to One-Tap Authentication Services in Cellular Networks (pg. 49)
- QPCASIN: A Quantum-Defended Privacy-Aware Preemptive Handover-Enabled Continuous Authentication in Space Information Networks (pg. 50)

### Autonomous Systems

- Development of Verification and Validation System for Security Assessment of ECU Components (pg. 70)

### Human Centric Security

- Security and Privacy in Augmented and Virtual Reality (AR/VR) (pg. 104)
- A Human Factors Perspective on Better Phishing Defenses (pg. 204)

### Network Security

- Formalizing Enterprise Firewall Management with Informal and Elastic Specifications (pg. 117)

### Prevention, Detection, and Response

- Quantitative Evaluation of Extensive Vulnerability Set Using Cost Benefit Analysis (pg. 134)
- Causality-Driven Mitigation of Cascading Failures in Distributed Systems (pg. 155)
- Testing and Detecting Software Upgrade Failures in Data-intensive Distributed Systems (pg. 156)

# Recently Added (Spring 2025)

### Artificial Intelligence and Machine Learning

- Malware Speaks! Deep Learning Based Assembly Code Processing for Detecting Evasive Cryptojacking (pg. 22)
- Artificial Intelligence for Music (pg. 29)
- Phishing Email Detection Through Machine Learning and Word Error Correction (pg. 33)

### Assured Identity and Privacy

- ABAUS: Active Bundle AUthentication Solution Based on SDN for Vehicular Networks (pg. 48)

### Autonomous Systems

- Development of Verification and Validation System for Security Assessment of ECU Components (pg. 70)

### Human Centric Security

- RUDOLF: An Efficient and Adaptive Defense Approach Against Website Fingerprinting Attacks Based on Soft Actor-Critic Algorithm (pg. 102)

### Network Security

- Explainable AI Methods for Enhancing AI-Based Network Intrusion Detection Systems (pg. 109)
- A Framework of High-Speed Network Protocol Fuzzing Based on Shared Memory (pg. 110)
- BioKA-ASVN: Biometric-Based Key Agreement Scheme for Air Smart Vehicular Networks Using Blockchain Service (pg. 111)
- Quantitative Evaluation of Extensive Vulnerability Set Using Cost Benefit Analysis (pg. 112)

### Prevention, Detection, and Response

- Scalable and Concurrent Targeted Search for Digital Forensics (pg. 142)

### Security Awareness, Education, and Training

- Crime Scene Surveying for IoT Investigations: National Training and Technical Assistance Program (pg. 158)

# Artificial Intelligence and Machine Learning

## Average Reward Reinforcement Learning

**PI: Vaneet Aggarwal**

### Overview

Most real world problems have infinite horizon average reward objectives, while this case has not been as well understood. The key reason is that the contraction operation that gives the key results in the discounted setup no longer holds. In our work, we aim to give the foundations of average reward reinforcement learning.

### Representative Publications

Bhrij Patel, Wesley A. Suttle, Alec Koppel, Vaneet Aggarwal, Brian M. Sadler, Amrit Singh Bedi, and Dinesh Manocha, "Global Optimality without Mixing Time Oracles in Average-reward RL via Multi-level Actor-Critic," in Proc. ICML, Jul 2024

Qinbo Bai, Washim Uddin Mondal, and Vaneet Aggarwal, "Regret Analysis of Policy Gradient Algorithm for Infinite Horizon Average Reward Markov Decision Processes," in Proc. AAAI, Feb 2024.

Qinbo Bai, Washim Uddin Mondal, and Vaneet Aggarwal, "Learning General Parameterized Policies for Infinite Horizon Average Reward Constrained MDPs via Primal-Dual Policy Gradient Algorithm," in Proc. Neurips, Dec 2024

Swetha Ganesh, Washim Uddin Mondal, Vaneet Aggarwal, "Variance-Reduced Policy Gradient Approaches for Infinite Horizon Average Reward Markov Decision Processes," arXiv, Apr 2024

Swetha Ganesh and Vaneet Aggarwal, "An Accelerated Multi-level Monte Carlo Approach for Average Reward Reinforcement Learning with General Policy Parametrization," arXiv, Jul 2024.

Project URL: https://web.ics.purdue.edu/~vaneet/publi_avrl.htm

# Intelligent Transportation Systems

**PI: Vaneet Aggarwal**

### Overview

The success of modern ride-sharing platforms crucially depends on the profit of the ride-sharing fleet operating companies, and how efficiently the resources are managed. Further, ride-sharing allows sharing costs and, hence, reduces the congestion and emission by making better use of vehicle capacities. The figure alongside depicts the improved performance of proposed strategy, DeepPool, for ride-sharing. The number of customers accepted are higher for same number of vehicles used and ride-sharing improves the costs, travel times, and number of customers served. The aspects are extended to joint pricing, matching, and dispatching problems. The approach is also used in freight management.

### Representative Publications

Xinwu Qian, Shuocheng Guo, and Vaneet Aggarwal, "DROP: Deep relocating option policy for optimal ride-hailing vehicle repositioning,"Transportation Research Part C, vol. 145, 103923, Dec 2022

Kaushik Manchella, Marina Haliem, Vaneet Aggarwal, and Bharat Bhargava, "PassGoodPool: Joint Passengers and Goods Fleet Management with Reinforcement Learning aided Pricing, Matching, and Route Planning,"IEEE Transactions on Intelligent Transportation Systems, vol. 23, no. 4, pp. 3866-3877, April 2022

Marina Haliem, Vaneet Aggarwal, and Bharat Bhargava, "AdaPool: A Diurnal-Adaptive Fleet Management Framework using Model-Free Deep Reinforcement Learning and Change Point Detection," IEEE Transactions on Intelligent Transportation Systems, vol. 23, no. 3, pp. 2471-2481, March 2022

Marina Haliem, Ganapathy Mani, Vaneet Aggarwal, and Bharat Bhargava, "A Distributed Model-Free Ride-Sharing Approach for Joint Matching, Pricing, and Dispatching using Deep Reinforcement Learning," IEEE Transactions on Intelligent Transportation Systems, vol. 22, no. 12, pp. 7931-7942, Dec. 2021.

Jiayu Chen, Abhishek K. Umrawal, Tian Lan, and Vaneet Aggarwal, "DeepFreight: A Model-free Deep-reinforcement-learning-based Algorithm for Multi-transfer Freight Delivery," in Proc. ICAPS, Aug 2021.

Marina Haliem, Vaneet Aggarwal, and Bharat Bhargava, "AdaPool: An Adaptive Model-Free Ride-Sharing Approach for Vehicle Dispatching using Deep Reinforcement Learning," in Proc. ACM Buildsys, Nov. 2020

Marina Haliem, Ganapathy Mani, Vaneet Aggarwal, and Bharat Bhargava, "Distributed Model-Free Ride-Sharing Algorithm with Pricing using Deep Reinforcement Learning," in Proc. ACM Computer Science in Cars Symposium (CSCS), Dec 2020

K. Manchella, A. K. Umrawal, and V. Aggarwal, "FlexPool: A Distributed Model-Free Deep Reinforcement Learning Algorithm for Joint Passengers and Goods Transportation," IEEE Transactions on Intelligent Transportation Systems, vol. 22, no. 4, pp. 2035-2047, April 2021.

Kaushik Manchella, Marina Haliem, Vaneet Aggarwal, and Bharat Bhargava, "A Distributed Delivery-Fleet Management Framework using Deep Reinforcement Learning and Dynamic Multi-Hop Routing," in Proc. Neurips Workshop on Machine Learning for Autonomous Driving, Dec 2020

Ashutosh Singh, Abubakr Alabbasi, and Vaneet Aggarwal, "A Distributed Model-Free Algorithm for Multi-hop Ride-sharing using Deep Reinforcement Learning," Accepted to IEEE Transactions on Intelligent Transportation Systems, May 2021.

Ashutosh Singh, Abubakr Alabbasi, and Vaneet Aggarwal, "A Reinforcement Learning Based Algorithm for Multi-hop Ride-sharing: Model-free Approach," in Proc. Neurips Workshop, Dec 2019.

A. Al-Abassi, A. Ghosh, and V. Aggarwal, "DeepPool: Distributed Model-free Algorithm for Ride-sharing using Deep Reinforcement Learning," IEEE Transactions on Intelligent Transportation Systems, vol. 20, no. 12, pp. 4714-4727, Dec. 2019 (Featured as ICAPS 2020 journal paper).

## Machine Learning Approach to Joint Multi-Agent Decision Making

**PI: Vaneet Aggarwal**

### Overview

Reinforcement Learning (RL) is being increasingly applied to optimize complex functions that may have a stochastic component. RL is extended to multi-agent systems to find policies to optimize systems that require agents to coordinate or to compete under the umbrella of Multi-Agent RL (MARL). A crucial factor in the success of RL is that the optimization problem is represented as the expected sum of rewards, which allows the use of backward induction for the solution. However, many real-world problems require a joint objective that is non-linear and dynamic programming cannot be applied directly. For example, in a resource allocation problem, one of the objective is to maximize long-term fairness among the users. This work addresses and formalizes the problem of joint objective optimization, where not only the sum of rewards of each agent but a function of the sum of rewards of each agent needs to be optimized

### Representative Publications

Mridul Agarwal and Vaneet Aggarwal, “A Reinforcement Learning Based Approach for Joint Multi-Agent Decision Making,” arXiv, Sept 2019.

Abubakr Al-Abbasi, Arnob Ghosh, and Vaneet Aggarwal, “DeepPool: Distributed Model-free Algorithm for Ride-sharing using Deep Reinforcement Learning,” Accepted to IEEE Transactions on Intelligent Transportation Systems, Jul 2019. https://arxiv.org/pdf/1903.03882.pdf

Ashutosh Singh, Abubakr Alabbasi, and Vaneet Aggarwal. “A Reinforcement Learning Based Algorithm for Multi-hop Ride-sharing: Model-free Approach.” In Proc. Neurips Workshop, 2019.

# Optimal Sample Complexity Guarantees for Discounted Reward Reinforcement Learning

**PI: Vaneet Aggarwal**

### Overview

Sample Complexity Guarantees for Discounted MDP in the model-free setups are far from lower bounds in many setups. In these works, we provide the state of art sample complexity results for policy-gradient based algorithms, actor-critic algorithms in the absence of constraints. Further, novel approaches are proposed to account for constraints in both tabular and parametrized setups.

### Representative Publications

Washim Uddin Mondal and Vaneet Aggarwal, "Improved Sample Complexity Analysis of Natural Policy Gradient Algorithm with General Parameterization for Infinite Horizon Discounted Reward Markov Decision Processes," in Proc. AISTATS, May 2024.

Mudit Gaur, Amrit Singh Bedi, Di Wang, and Vaneet Aggarwal, "Closing the Gap: Achieving Global Convergence (Last Iterate) of Actor-Critic under Markovian Sampling with Neural Network Parametrization," in Proc. ICML, Jul 2024 (Spotlight Paper)

Washim Uddin Mondal and Vaneet Aggarwal, "Sample-Efficient Constrained Reinforcement Learning with General Parameterization," in Proc. Neurips, Dec 2024.

Washim Uddin Mondal and Vaneet Aggarwal, "Last-Iterate Convergence of General Parameterized Policies in Constrained MDPs," arXiv, Aug 2024.

Project URL: https://web.ics.purdue.edu/~vaneet/publi_dirl.htm

# Quantum Machine Learning

**PI: Vaneet Aggarwal**

### Overview

Quantum Machine Learning (QML) is an emerging research area advocating the use of quantum computing for advancement in machine learning. In this project, we aim to demonstrate the speedup in machine learning tasks due to the use of quantum computing

### Representative Publications

Bhargav Ganguly, Yulian Wu, Di Wang, and Vaneet Aggarwal, "Quantum Computing Provides Exponential Regret Improvement in Episodic Reinforcement Learning," Feb 2023.

Dheeraj Peddireddy, Utkarsh Priyam, and Vaneet Aggarwal, "Noisy Tensor Ring approximation for computing gradients of Variational Quantum Eigensolver for Combinatorial Optimization," Feb 2023

Yulian Wu, Chaowen Guan, Vaneet Aggarwal, and Di Wang, "Quantum Heavy-tailed Bandits," Jan 2023.

Dheeraj Peddireddy, Vipul Bansal, and Vaneet Aggarwal, "Classical simulation of variational quantum classifiers using tensor rings," Applied Soft Computing, Volume 141, 110308, July 2023.

Debanjan Konar, Aditya Das Sarma, Soham Bhandary, Siddhartha Bhattacharyya, Attila Cangia, and Vaneet Aggarwal, "A Shallow Hybrid Classical-Quantum Spiking Feedforward Neural Network for Noise-Robust Image Classification," Applied Soft Computing, vol. 136, paper 110099, Mar 2023

Mohammad Ali Javidian, Vaneet Aggarwal, and Zubin Jacob, "Quantum Causal Inference in the Presence of Hidden Common Causes: an Entropic Approach," Physical Review A, 106, 062425, Dec 2022.

Mohammad Ali Javidian, Vaneet Aggarwal, and Zubin Jacob, "Learning Circular Hidden Quantum Markov Models: A Tensor Network Approach," arXiv, Oct 2021

Dheeraj Peddireddy, Vipul Bansal, Zubin Jacob, and Vaneet Aggarwal, "Tensor Ring Parametrized Variational Quantum Circuits for Large Scale Quantum Machine Learning," in Proc. Neurips Workshop on Quantum Tensor Networks in Machine Learning, Dec. 2021.

Mohammad Ali Javidian, Vaneet Aggarwal, and Zubin Jacob, "Tensor Rings for Learning Circular Hidden Markov Models," in Proc. Neurips Workshop on Quantum Tensor Networks in Machine Learning, Dec. 2021.

Project URL: https://web.ics.purdue.edu/~vaneet/publi_quantum.htm

## Reinforcement Learning for Network Resilience

**PI: Vaneet Aggarwal**

*Funding Source: Cisco, Meta, Inc.*

**Overview**

With the rapid growth of backbone networks and data center networks, ensuring network robustness and resilience has become a key challenge in network design. In this project, we use multi-agent reinforcement learning approaches for content placement, satisfying different peak and average constraint requirements in decision making.

**Representative Publications**

Chang-Lin Chen, Hanhan Zhou, Jiayu Chen, Mohammad Pedramfar, Vaneet Aggarwal, Tian Lan, Zheqing Zhu, Chi Zhou, Tim Gasser, Pol Mauri Ruiz, Vijay Menon, Neeraj Kumar, and Hongbo Dong, "Two-tiered Online Optimization of Region-wide Datacenter Resource Allocation via Deep Reinforcement Learning," Jun 2023.

Chenyi Liu, Vaneet Aggarwal, Tian Lan, Nan Geng, Yuan Yang, and Mingwei Xu, "Machine Learning for Robust Network Design: A New Perspective," Accepted to IEEE Communications Magazine, May 2023.

Chenyi Liu, Vaneet Aggarwal, Tian Lan, Nan Geng, Yuan Yang, Mingwei Xu, and Qing Li, "FERN:

Leveraging Graph Attention Networks for Failure Evaluation and Robust Network Design," Accepted to IEEE/ACM Transactions on Networking, Aug 2023.

Nan Geng, Qinbo Bai, Chenyi Liu, Tian Lan, Vaneet Aggarwal, Yuan Yang, and Mingwei Xu, "A Reinforcement Learning Framework for Vehicular Network Routing Under Peak and Average Constraints," IEEE Transactions on Vehicular Technology (TVT), vol. 72, no. 5, pp. 6753-6764, May 2023.

Chenyi Liu, Nan Geng, Vaneet Aggarwal, Tian Lan, Yuan Yang and Mingwei Xu, "CMIX: Deep Multi-agent Reinforcement Learning with Peak and Average Constraints" in Proc. ECML, Sep 2021

Nan Geng, Tian Lan, Vaneet Aggarwal, Yuan Yang, and Mingwei Xu, "A Multi-agent Reinforcement Learning Perspective on Distributed Traffic Engineering," in Proc. IEEE International Conference on Network Protocols (ICNP), Oct 2020.

Yimeng Wang, Yongbo Li, Tian Lan, and Vaneet Aggarwal, "DeepChunk: Deep Q-Learning for Chunk-based Caching in Data Processing Networks," IEEE Transactions on Cognitive Communications and Networking, Special Issue on Deep Reinforcement Learning for Future Wireless Communication Networks, vol. 5, no. 4, pp. 1034-1045, Dec. 2019.

Ramkumar Raghu, Pratheek Upadhyaya, Mahadesh Panju, Vaneet Aggarwal, and Vinod Sharma, "Deep Reinforcement Learning Based Power control for Wireless Multicast Systems," in Proc. Allerton, Oct 2019.

Yimeng Wang, Yongbo Li, Vaneet Aggarwal, and Tian Lan, "Deep Q-Learning for Chunk-based Caching in Data Processing Networks," in Proc. Allerton, Oct 2019.

Project URL: https://web.ics.purdue.edu/~vaneet/publi_ml_net.htm

# Submodular Bandits

**PI: Vaneet Aggarwal**

### Overview

We investigate the problem of stochastic, combinatorial multi-armed bandits where the learner only has access to bandit feedback and the reward function can be non-linear. We provide a general framework for adapting discrete offline approximation algorithms into sub-linear regret methods that only require bandit feedback. The framework only requires the offline algorithms to be robust to small errors in function evaluation. The adaptation procedure does not even require explicit knowledge of the offline approximation algorithm -- the offline algorithm can be used as black box subroutine. Such approaches are useful in wide variety of setups including social influence maximization, revenue management, market design, product ranking optimization in online platforms, and reserve price optimization in auctions. Further, many of these problems have continuous combinatorial actions, which is also being explored.

### Representative Publications

Mohammad Pedramfar, Christopher John Quinn, and Vaneet Aggarwal, "A Unified Approach for Maximizing Continuous DR-submodular Functions," May 2023

Mohammad Pedramfar and Vaneet Aggarwal, "Stochastic Submodular Bandits with Delayed Composite Anonymous Bandit Feedback," Mar 2023.

Guanyu Nie, Yididiya Y Nadew, Yanhui Zhu, Vaneet Aggarwal, and Christopher John Quinn, "A Framework for Adapting Offline Algorithms to Solve Combinatorial Multi-Armed Bandit Problems with Bandit Feedback,"in Proc. ICML, Jul 2023.

Fares Fourati, Vaneet Aggarwal, Christopher Quinn, and Mohamed-Slim Alouini, "Randomized Greedy Learning for Non-monotone Stochastic Submodular Maximization Under Full-bandit Feedback," in Proc. AISTATS, Apr 2023.

Guanyu Nie, Mridul Agarwal, Abhishek Kumar Umrawal, Vaneet Aggarwal, Christopher John Quinn, "An Explore-then-Commit Algorithm for Submodular Maximization Under Full-bandit Feedback," in Proc. UAI, Aug 2022.

Mridul Agarwal, Vaneet Aggarwal, Christopher J. Quinn, and Abhishek Umrawal, "Stochastic Top K-Subset Bandits with Linear Space and Non-Linear Feedback with Applications to Social Influence Maximization," ACM/IMS Transactions on Data Science, vol. 2, issue 4, Article 38, Nov 2021.

Mridul Agarwal, Vaneet Aggarwal, Christopher J. Quinn, and Abhishek Umrawal, "DART: aDaptive Accept RejecT for non-linear top-K subset identification," in Proc. AAAI, Feb 2021 (21% acceptance rate, 1692/7911).

Mridul Agarwal, Vaneet Aggarwal, Christopher J. Quinn, and Abhishek Umrawal, "Stochastic Combinatorial Bandits with Linear Space and Non-Linear Feedback," in Proc. ALT, Mar 2021 (PMLR 132:306-339, 2021.) (29.3% acceptance rate, 46/157).

Project URL: https://web.ics.purdue.edu/~vaneet/publi_bandits.htm

## Temporal Abstractions in Multi-Agent Reinforcement Learning

**PI: Vaneet Aggarwal**

### Overview

Covering option discovery has been developed to improve the exploration of reinforcement learning in single-agent scenarios with sparse reward signals, through connecting the most distant states in the embedding space provided by the Fiedler vector of the state transition graph. However, these option discovery methods cannot be directly extended to multi-agent scenarios, since the joint state space grows exponentially with the number of agents in the system. In order to alleviate this problem, we design efficient approaches to make multi-agent deep covering options scalable.

The proposed multi-agent exploration approaches can be used for learning how multiple robots can pick up the object together, coordinate to move across doors, without explosion in complexity. Scalable algorithms are provided.

### Representative Publications

Jiayu Chen, Jingdi Chen, Tian Lan, and Vaneet Aggarwal, "Multi-agent Covering Option Discovery based

on Kronecker Product of Factor Graphs," Accepted IEEE TAI, 2022.

Jiayu Chen, Jingdi Chen, Tian Lan, and Vaneet Aggarwal, "Multi-agent Covering Option Discovery through Kronecker Product of Factor Graphs," in Proc. AAMAS, May 2022

Jiayu Chen, Tian Lan, Vaneet Aggarwal, "Hierarchical Adversarial Inverse Reinforcement Learning for Robotic Manipulation," in Proc. IEEE International Conference on Robotics and Automation (ICRA), May 2023.

Jiayu Chen, Jingdi Chen, Tian Lan, and Vaneet Aggarwal, "Scalable Multi-agent Covering Option Discovery based on Kronecker Graphs," in Proc. Neurips, Dec 2022.

Jiayu Chen, Vaneet Aggarwal, and Tian Lan, "ODPP: A Unified Algorithm Framework for Unsupervised Option Discovery based on Determinantal Point Process," Dec 2022.

Jiayu Chen, Dipesh Tamboli, Tian Lan, and Vaneet Aggarwal, "Multi-task Hierarchical Adversarial Inverse Reinforcement Learning," in Proc. ICML, Jul 2023.

## Malware Speaks! Deep Learning Based Assembly Code Processing for Detecting Evasive Cryptojacking

**PI: Bharat Bhargava**

**Current Students: Ganapathy Mani; Myeongsu Kim; Pelin Angin; Ayça Deniz; Vikram Pasumarti**

### Overview

The increasing prevalence of blockchain-based cryptocurrencies as a payment instrument in the past decade and the rewards earned by the cryptominers has resulted in a new class of cyber attacks, cryptojacking , which involves unauthorized mining of cryptocurrencies on someone's system. Spotting cryptojacking is difficult in many cases, since the relevant software tries to disguise its presence to evade detection, by mimicking benign software such as compression applications by performing similar bitwise, cryptographic, and encryption operations. In this paper, we propose the processing of assembly code—a fundamental and platform-independent programming language—as a natural language using deep learning for profiling applications, which we call De ep Code Pro filer (DeCode Pro). Our proposed solution leverages the immutable step of any cyber attack: the deployment of instructions in system memory to carry out the attack. Through extensive experimentation with different neural network architectures in the profiling stage, we show that DeCode Pro is highly effective in the detection of evasive cryptojacking attacks and achieves low false positive and false negative rates. We also show that the model achieves high classification accuracy even with limited training data, which can considerably reduce the computing resources required for training and retraining the deep learning model.

### Representative Publications

IEEE Transactions on Dependable and Secure Computing ( Volume: 21, Issue: 4, July-Aug. 2024)

Project URL: https://ieeexplore.ieee.org/document/10226222

# DARPA Project GNOME with USC

**PI: Bharat Bhargava**

**Current Students: Marina Wagdy Wadea Haliem, KMA Solaiman, Alina V Nesen, Mijanur Palash, Trevor A Bonjour, Shafkat Islam, MD Shamsul Kaonain**

### Overview

GNOME is a computational platform for generating, studying and analyzing novelty in the context of open-world Artificial Intelligence (AI). Open-world intuitively means that the world is not fully specified, known or parameterized in advance. The real world, as seen from the human perspective, is the best example. In the real world, human beings are constantly having to adapt to novelty, both in the long and short term. In contrast, AI systems are far less versatile, despite having achieved near human-level performance in specific, highly scoped problem areas such as face recognition. As a platform, GNOME allows researchers to simulate and 'play with' novelty as a first-class citizen. GNOME also serves as a testbed for evaluating novelty-adaptive AI agents in strategic gameplaying environments.

GNOME is funded under the DARPA SAIL-ON (Science of Artificial Intelligence and Learning for Open-world Novelty) program. While current AI systems excel at tasks defined by 'rigid' rules, they are not very good at adapting to changing conditions commonly faced by people in the real world: from driving in fluctuating weather, walking on uneven terrain, playing games with slightly different rules (e.g., chess under a time constraint) and, in general, modifying the internal model of the world when an assumption turns out to be faulty.

Source: https://usc-isi-i2.github.io/gnome/

Project URL: https://usc-isi-i2.github.io/gnome/

# Security and Privacy of Large Language Models (LLMs)

**PI: Berkay Celik**

### Overview

This project explores the security and privacy issues associated with Large Language Models (LLMs). Research in this area assesses the capability of LLMs to offer correct security advice and refute common misconceptions. It also involves developing new methods to evaluate "jailbreak" attacks that aim to circumvent the safety mechanisms of these models.

Project URL: https://beerkay.github.io/

# FAI: Identifying, measuring, and mitigating fairness issues in AI

**PIs: Chris Clifton, Chris Yeomans, Lindsay Weinberg, Murat Kantarcioglu (University of Texas at Dallas), Blase Ur (University of Chicago)**

**Current Students: Rakin Haider, Ryan Van Nood**

*Funding Source: Amazon, National Science Foundation (NSF)*

### Overview

Bias and Discrimination in Artificial Intelligence (AI) has been receiving increasing attention. Unfortunately, the positive concept Fair AI is difficult to define. For example, it is hard to distinguish between (desired) personalization and (undesired) bias. These differences often depend on context, such as the use of gender or ethnicity in making a medical diagnosis vs. using the same attributes in determining if insurance should cover a medical procedure. This is particularly difficult as AI systems are used in new contexts, enabling products and services that have not been seen before and for which societal concepts of fairness are not yet established. This multidisciplinary project will construct a framework and taxonomy for understanding fairness in societal contexts. Human-computer interaction methods will be developed to learn perceptions of fairness based on human interaction with AI systems. Automated methods will be developed to relate these perceptions to the framework, enabling developers (and eventually automated AI systems) to respond to and correct issues perceived by users of the systems.

This exploratory project will develop a taxonomy incorporating concepts of Aristotelian fairness (distributive vs. corrective justice) and Rawlsian fairness (equality of rights and opportunities). A formal literature survey will be used to establish a framework for societal contexts of fairness and how they relate to the Taxonomy. Experiments with perceptions of models both in isolation and in comparison will be used to evaluate situations where people perceive AI systems as fair or unfair. Tools will be developed to identify and explain fairness issues in terms of the taxonomy, based on the elicited perceptions and societal context of the system. While beyond the scope of this project, the outcome of these tools could potentially be used to automatically adjust AI systems to reduce unfairness.

# Exploring the Process, Challenges, and Effective Practices of Reproducing and Adapting Machine Learning Models

**PIs: Jamie Davis (Purdue), Yung-Hsiang Lu (Purdue), George K Thiruvathukal (Loyola University Chicago)**

**Current Students: Wenxin Jiang, PhD student, ~20 Purdue undergraduate students from several majors, through the VIP program**

### Overview

Complex computing systems increasingly rely on components derived from machine learning and data science. Machine learning and data science techniques have been adopted across most business enterprises. These techniques include the development of machine learning models, and the use of analysis pipelines to automatically and repeatedly process batches of data. Engineering these models, and

reproducing and extending analysis pipelines, are critical aspects of modern computing. It is also hard! Getting them right is a major challenge. This line of research has many applications in the 21st century. For example, enabling reproducible, high-quality machine learning and computer vision software will support efforts for national security and national defense.

In this project, we are studying the process by which "research prototype" ML models are translated into high-quality engineering artifacts suitable for use in the original or new contexts. This process has aspects of scientific replication, and of technology transfer. We are applying software engineering research methods to characterize the process, identify pitfalls, and document effective practices. We are studying general practices, as well as specific practices in low-power/embedded contexts.

This work involves a partnership with Google.

### Representative Publications

Exemplars for Machine Learning: Towards Software Engineering & Reproducibility. Vivek, Chinnakotla, Banna, Vegesana, Yan, Davis, Lu, and Thiruvathukal. SIAM Conference on Computational Science and Engineering (CSE'20) 2020.

Project URL: https://engineering.purdue.edu/VIP/teams/tensorflow

## Reusing Deep Learning Models: Challenges and Directions in Trustworthy Software Engineering

**PIs: Jamie Davis, Yung-Hsiang Lu (Purdue), George K Thiruvathukal (Loyola University Chicago)**

**Current Students: Purvish Jajal, PhD student, Nick Eliopoulos, PhD student, Parth Patil, MSc student, ~20 Purdue undergraduate students from several majors, through the VIP program**

*Funding Source: Cisco, Google, Inc. (partial funding), National Science Foundation (NSF)*

### Overview

The development and training of deep learning models have become increasingly costly and complex. Consequently, software engineers are adopting pre-trained models (PTMs) for their downstream applications. The content, dynamics, and effective use of the PTM supply chain remain largely unexplored. This project seeks to characterize the associated engineering processes and artifacts, in order to identify and mitigate failure modes. We then develop tools, e.g. automation, to optimize the secure reuse of PTMs. We integrate research methods from human factors, mining software repositories, and machine learning. We are focused on major model registries such as HuggingFace and PyTorchHub, and interoperability infrastructure such as Pickle (pickle deserialization attacks) and the Open Neural Network eXchange (ONNX -- converter degradation challenges).

This project's goal is to accelerate secure software engineering work with PTMs. For example, we have built:

- The first typosquatting detector targeting pre-trained models, also achieving state-of-art detection

performance on other registries such as NPM and PyPI
- The first secure loader for Pickle-based ML models
- The first analysis tool to detect incorrectly named ML models (eg backdoors)

**Representative Publications**

An Empirical Study of Pre-Trained Model Reuse in the Hugging Face Deep Learning Model Registry.
Jiang, Synovic, Hyatt, Schorlemmer, Sethi, Lu, Thiruvathukal, and Davis.
Proceedings of the ACM/IEEE 45th International Conference on Software Engineering (ICSE) 2023.

Reusing Deep Learning Models: Challenges and Directions in Software Engineering.
Davis, Jajal, Jiang, Schorlemmer, Synovic, and Thiruvathukal.
Proceedings of the IEEE John Vincent Atanasoff Symposium on Modern Computing (JVA'23) 2023.

Analysis of Failures and Risks in Deep Learning Model Converters: A Case Study in the ONNX Ecosystem.
Jajal, Jiang, Tewari, Woo, Lu, Thiruvathukal, and Davis.
arXiv 2023.

An Empirical Study of Artifacts and Security Practices in the Pre-trained Model Supply Chain.
Jiang, Synovic, Sethi, Indarapu, Hyatt, Schorlemmer, Thiruvathukal, and Davis.
Proceedings of the 1st ACM Workshop on Software Supply Chain Offensive Research and Ecosystem Defenses (SCORED) 2022.

Discrepancies among Pre-trained Deep Neural Networks: A New Threat to Model Zoo Reliability.
Montes, Peerapatanapokin, Schultz, Guo, Jiang, and Davis.
Proceedings of the 30th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering: Ideas, Visions, and Reflections track (ESEC/FSE-IVR) 2022.

Exploring Naming Conventions (and Defects) of Pre-trained Deep Learning Models in Hugging Face and Other Model Hubs.
Jiang, Cheung, Thiruvathukal, and Davis. arXiv 2023.

AgentHub: A Research Agenda for Agent Sharing Infrastructure E Pautsch, T Singla, W Jiang, H Peng, B Hassanshahi, K Läufer, G.K. Thiruvathukal, J.C. Davis

Software Dependencies 2.0: An Empirical Study of Reuse and Integration of Pre-Trained Models in Open-Source Projects
J Yasmin, W Jiang, JC Davis, Y Tian

PickleBall: Secure Deserialization of Pickle-based Machine Learning Models
AD Kellas, N Christou, W Jiang, P Li, L Simon, Y David, VP Kemerlis, - CCS 2025

Advancing Jailbreak Strategies: A Hybrid Approach to Exploiting LLM Vulnerabilities and Bypassing Modern Defenses - M Ahmed, M Abdelmouty, M Kim, G Kandula, A Park, JC Davis - SecDev'25-Poster

AI Safety in the Eyes of the Downstream Developer: A First Look at Concerns, Practices, and Challenges - H Gao, M Zahedi, W Jiang, HY Lin, J Davis, C Treude

ConfuGuard: Using Metadata to Detect Active and Stealthy Package Confusion Attacks Accurately and at Scale - W Jiang, B Çakar, M Lysenko, JC Davis

Recommending Pre-Trained Models for IoT Devices - Parth V Patil, Wenxin Jiang, Huiyun Peng, Daniel Lugo, Kelechi G Kalu, Josh LeBlanc, Lawrence Smith, Hyeonwoo Heo, Nathanael Aou, James C Davis - SERP4IoT'25

## Using GenAI in Software Engineering Education: Efficiency and Cybersecurity

PI: Jamie Davis, Dr. Kirsten Davis (Engineering Education)

*Funding Source: National Science Foundation (NSF)*

### Overview

Software enginering students need to learn how to incorporate GenAI to accelerate their work, without compromising on its quality. In ECE courses we are incorporating GenAI learning modules and coaching students on its use.

### Representative Publications

An Exploratory Study on Upper-Level Computing Students' Use of Large Language Models as Tools in a Semester-Long Project. Tanay, Arinze, Joshi, Davis, and Davis.
Annual Conference of the American Society for Engineering Education (ASEE) 2024.

## Pose and Range Estimation from Monocular Monstatic Single Resolved Images

### PI: Carolin Frueh

### Overview

The range and attitude state of an unknown spacecraft is estimated using a single resolved EO image using a combination of machine learning and traditional techniques. Fully autonomous.

# Automatic Discovery of Real-Time Visual Data on the Internet

**PI: Yung-Hsiang Lu** **Current Students: Ryan Merrill Dailey**

### Overview

Many network cameras are connected to the Internet provide real-time visual data (image or video) all around the world. As computer vision technologies become widely available, it is now feasible analyzing the data streams and responding to situations as they are developing. Network cameras may be observing severe weather and computer vision may detect people needing help. Some network cameras may monitor sensitive areas and computer vision may detect unauthorized access. In the era of machine learning, realistic training data plays crucial roles to success. The diverse data from network cameras can be used to improve and evaluate learning-based computer vision. Currently there is no repository storing the information about these network cameras. This project is building an automatic method to discover network cameras on the Internet.

### Representative Publications

"Creating the World's Largest Real-Time Camera Network", Imaging and Multimedia Analytics in a Web and Mobile World 2017

Project URL: https://www.cam2project.net/

# NSF CRII: RI: A Study of Rank-based Decomposable Losses for Machine Learning

**PI: Shu Hu** **Current Students**: Li Lin, Santosh

### Overview

This project will be conducted in two interrelated thrusts. The first thrust explores a novel and general rank-based aggregate loss for supervised learning. The focus will encompass efficient algorithms that can optimize this loss with guaranteed convergence, along with streamlined techniques to determine relevant hyperparameters. The developed loss will be connected with distributionally robust optimization to gain insights into its sample-level robustness and the development of new types of rank-based aggregate losses. Additionally, theoretical guarantees will be established for rank-based aggregate losses, including classification calibration, classification consistency, and generalization properties. The second thrust aims to study a general formulation of rank-based individual loss with theoretical analysis, bolstering label-level robustness in multi-class and multi-label learning scenarios. Furthermore, the use of rank-based individual loss will be expanded to tackle fairness learning challenges and investigate the resilience of models trained with this loss against adversarial threats, including verification and defense mechanisms.

Project URL: https://www.nsf.gov/awardsearch/showAward?AWD_ID=2434967

## Active Learning using Loss Geometry

**PI: Rajiv Khanna**

**Overview**
This work targets "atypical" points the model misclassifies with high confidence, shifting early rounds toward the right kind of hard examples. Aggregate tables show top average ranksover budgets and open-set ratios across CIFAR-10/100 and TinyImageNet.

## Machine Unlearning

**PI: Rajiv Khanna**

**Overview**
We analyze unlearning when a model must forget a designated subset while retaining utility on the rest, characterizing howsharpness-oriented loss landscape via optimization interacts with forget/retain signals and when it helps or hurts—providing theory-guided trade-offs and diagnostics.

## Scalable Data Efficient Learning

**PI: Rajiv Khanna**

Overview
We frame coreset selection (ie data subset selection) via posterior sampling so the subset's induced loss landscape better matches the full-data landscape, yielding robustness under label corruption and small-budget regimes and improving time-to-accuracy versus state of the art baselines. We report 20–200% time-to-best-accuracy gains and lower memory than current state-of-the-art across vision/NLP benchmarks,

## Artificial Intelligence for Music

**PIs: Yung-Hsiang Lu, Kristen Yeon-Ji Yun**

**Current Students: Benjamin Shiue-Hal Chou**

**Overview**
Beginner musicians often struggle to identify specific errors in their performances, such as playing incorrect notes or rhythms. There are two limitations in existing tools for music error detection: (1) Existing approaches rely on automatic alignment; therefore, they are prone to errors caused by small deviations between alignment targets.; (2) There is a lack of sufficient data to train music error detection models, resulting in over-reliance on heuristics. To address (1), we propose a novel transformer model, Polytune, that takes audio inputs and outputs annotated music scores. This model can be trained end-to-end to implicitly align and compare performance audio with music scores through latent space representations. To address (2), we present a novel data generation technique capable of creating large-

scale synthetic music error datasets. Our approach achieves a 64.1% average Error Detection F1 score, improving upon prior work by 40 percentage points across 14 instruments. Additionally, compared with existing transcription methods repurposed for music error detection, our model can handle multiple instruments.

### Representative Publications

Benjamin Shiue-Hal Chou, Purvish Jajal, Nicholas John Eliopoulos, Tim Nadolsky, Cheng-Yun Yang, Nikita Ravi, James C. Davis, Kristen Yeon-Ji Yun, Yung-Hsiang Lu, Detecting Music Performance Errors with Transformers, 2025 Annual Conference of Artificial Intelligence

Project URL: https://ai4musicians.org/

## Computer Vision using Contextual Information

**PI: Yung-Hsiang Lu**

**Current Students: Caleb Tung (tung3@purdue.edu)**

### Overview

Most computer vision solutions are designed to consider only pixels in images or videos. These solutions have to infer the environment about when and where the pixels are acquired. Modern cameras are often equipped positioning capabilities and can embed information about time and locations in images or videos. Such information can offer the context of the pixels. Contextual information may improve computer vision technologies in many ways. For example, many vehicles are expected during rush hours (time) in the downtown of a city (location). In contrast, few people are expected during semester breaks (time) on a university campus (location). Contextual information can be used to improve computer vision in many ways: First, the information can evaluate correctness: for example, an elephant is not expected in a city downtown. Second, such information may help to trim deep neural networks by removing impossible scenarios; for example, a traffic camera does not need the ability to recognize an elephant. Third, smaller neural networks may be deployed in edge devices that can perform computer vision without sending pixels through networks.

### Representative Publications

See the World through Network Cameras (accepted by IEEE Computer for publication)

# Dynamic Optimization Using Object Detection in Bandwidth Constrained Automatic Multi-camera Networks

**PI: Yung-Hsiang Lu**

**Current Students: Haobo Wang, Master Degree, Purdue University**

### Overview

Large-scale multi-camera networks are getting an increasing demand for visual surveillance and multi-target tracking. However, there has been a challenge for the multi-camera surveillance system due to the limitation of bandwidth. Most recent researches on such system design consider either unlimited network resources, or under the worst case assumption which might lose information for uncovered areas. In this paper, we present a novel system to optimize the information obtained from multi-camera visual data given bandwidth limitations, which utilizes both the camera placement maximization and object detection algorithm. Our proposed method can define the priority of monitored cameras based on real-time video context and request different specs of data from individual cameras dynamically. We also suggest a cost-efficient, self-controllable multi-camera network emulation testbed, using minimega made by Sandia National Lab.

# Low-Power Computer Vision

**PI: Yung-Hsiang Lu** **Current Students: Abhinav Goel**

### Overview

Learning-based computer vision requires vast amounts of computation passing information through many layers of neural networks. This project investigates how to improve energy efficiency of computer vision so that it can run on embedded systems, thus at the cameras, without transmitting video streams through networks. State-of-the-art computer vision aims to create general-purpose solutions that can recognize hundreds of classes of objects. In constract, this project creates much smaller neural networks that specialize on distingushing only a few types of objects. If more types need to be distinguished, this project creates hierarchies of neural networks and allow early termination when detected objects are not of interest. This project has already demonstrated significant reduction in energy consumption with negligible loss of accuracy.

### Representative Publications

"Low-Power Computer Vision: Status, Challenges, Opportunities", IEEE Journal on Emerging and Selected Topics in Circuits and Systems. Volume: 9 , Issue: 2 , June 2019

"Low-power image recognition", Nature Machine Learning Vol 1, April 2019,
"Low-Power Image Recognition Challenge", AI Magazine Vol 39 No 2, Summer 2018

Project URL: https://www.cam2project.net/

# Securing Data Privacy at the Edge Using Trusted Execution Environment

**PI: Yung-Hsiang Lu**

### Overview

Data privacy on edge devices has become a significant concern, given the widespread use of mobile phones. There is a potential risk of sensitive personal data being leaked through an insecure network domain or a vulnerable system update. Previous efforts to address data privacy at the edge have either introduced substantial computational overhead or led to accuracy degradation in the original task.

This project proposes the utilization of Trusted Execution Environment (TEE), a hardware feature that establishes an isolated environment on edge devices for secure deep neural network (DNN) model inference without compromising accuracy. Additionally, we have developed compression methods for DNNs to mitigate the computational overhead associated with running DNN models in TEE at the edge.

# Development of a Reliable Method for General Aviation Flight Phase Identification

**PI: John Mott**

### Overview

Aircraft operations statistics have typically received significant attention from U.S. airport owners and operators and state, local, and federal agencies. Accurate operational data is beneficial in assessing airports' performance efficiency and impact on the environment, but operational statistics at nontowered general aviation airports are, for the most part, limited or not available. However, the increasing availability and economy of capturing and processing Automatic Dependent Surveillance-Broadcast (ADS-B) data shows promise for improving accessibility to a wide variety of information about the aircraft operating in the vicinity of these airports. Using machine learning technology, specific operational details can be decoded from ADS-B data. This paper aims to develop a reliable and economical method for general aviation aircraft flight phase identification, thereby leading to improved noise and emissions models, which are foundational to addressing many public concerns related to airports.

### Representative Publications

Q. Zhang, J. H. Mott, M. E. Johnson and J. A. Springer, "Development of a Reliable Method for General Aviation Flight Phase Identification," in *IEEE Transactions on Intelligent Transportation Systems,* doi: 10.1109/TITS.2021.3106774.

Zhang, Q., & Mott, J. H. (2022). An improved framework for classification of flight phases of general aviation aircraft. Transportation *Research Record: Journal of the Transportation Research Board*.

# AI/Control for Intrusion Detection in Constrained Embedded Systems

**PIs: Shaoshuai Mou, Ding Zhao**

**Current Students: Paulo Heredia, Taashi Kapoor**

*Funding Source: Rolls-Royce*

### Overview

Embedded systems (ES) are key elements of real-world applications, which can be leveraged by cyber attackers to tamper the whole system. Enhancing the security of ES is of increasing interest especially when attackers are equipped with recent advances in artificial intelligence (AI) . Detection of intrusion is particularly challenging when it comes to constrained ES, such as the Engine Control Unit (ECU) in an aircraft, which works in harsh environments with high temperature and a lot of noise, and has limited processors against cyber-attacks. Although it is rare for constrained ES to be under attack thanks to its closed operational environment, cybersecurity of constrained ES is usually of extremely high value due to their safety critical applications such as ECU. As the first layer of all cyber-defense techniques, Intrusion Detection (ID) could detect cyber-attacks in an early stage, alert the ES to take early actions to mitigate the cyber-attacks and prevent further damages. Since signature-based approaches require a predefined model of attacks for intrusion detection and are not effective in detecting new attacks, research efforts have recently given to anomaly-based ID, in which intrusion is detected by comparing the observed behaviors according to the known normal behaviors. Anomaly-based ID has become a more promising tool with the advance of AI and machine learning (ML). Such AI-based IDs are able to process large volumes of data, do not require an exact knowledge of the system model, and improve performance with experience. Meanwhile, these techniques are also far from ideal, especially suffering from heavy computational load and high positive alarm rate. We also note that classical theories in control are well-developed and have played a key role in maintaining system's performance under noise and uncertainty, and detecting systems' faults online. Recognition of this has motivated us in this proposal to integrate recent advances in AI and ML with well-developed techniques in control and optimizations to develop innovative solutions to intrusion detection for constrained embedded systems.

# Phishing Email Detection Through Machine Learning and Word Error Correction

**PI: Quamar Niyaz**

### Overview

Phishing is one of the most prevalent and effective fraudulent activities on the Internet. Numerous machine learning (ML)-based models have been developed to detect phishing emails using publicly available datasets (e.g., Nazario, Millersmile). These email datasets have poor grammar structure or incorrect word usage, which ML models often learn as key distinguishing features. With the advent of large language models (LLMs), the grammatical quality and structure of phishing emails have significantly improved, making them appear more legitimate. As a result, traditional ML models that rely on

grammatical cues may become less effective in identifying phishing emails. To address this challenge, we explore the following research question: Can an ML-based phishing detection model, enhanced with word correction and splitting techniques, effectively identify phishing emails? To investigate this, we develop a phishing detection system that integrates misspelled word correction and combined-word splitting during the data preprocessing stage. The system leverages state-of-the-art natural language processing (NLP) techniques to enhance detection accuracy. Additionally, to improve model robustness, we utilize datasets from diverse sources and time periods for training and deployment.

### Representative Publications

Deeksha Kulal, Leul Shiferaw and Quamar Niyaz, "Phishing Email Detection Through Machine Learning and Word Error Correction," 2025 17th International Conference on COMmunication Systems and NETworks (COMSNETS), Bengaluru, India, 2025, pp. 1299-1304, doi: 10.1109/COMSNETS63942.2025.10885558.

## Data Preprocessing and ML Model Fairness

### PI: Romila Pradhan

### Current Students: Ekta, Sathvika Kotha

### Overview

The success of machine learning techniques in widespread applications has taught us that with respect to accuracy, the more data, the better the model. However, for fairness, data quality is perhaps more important than quantity. Before being fed into an ML model, training data undergoes a number of preprocessing steps. Existing studies have considered the impact of data preprocessing on the accuracy of ML model tasks. However, the impact of preprocessing on the fairness of the downstream model has neither been studied nor well-understood. In this project, we conduct a systematic study of how data quality issues and data preprocessing steps impact model fairness. Furthermore, we develop solutions for improving individual data preprocessing steps that would improve downstream model fairness.

## Explainable AI: Foundations, Applications, Opportunities for Data Management Research

### PI: Romila Pradhan

### Overview

Algorithmic decision-making systems are successfully being adopted in a wide range of domains for diverse tasks. While the potential benefits of algorithmic decision-making are many, the importance of trusting these systems has only recently attracted attention. There is growing concern that these systems are complex, opaque and non-intuitive, and hence are difficult to trust. There has been a recent resurgence of interest in explainable artificial intelligence (XAI) that aims to reduce the opacity of a model by explaining its behavior, its predictions or both, thus allowing humans to scrutinize and trust the model. A host of technical advances have been made and several explanation methods have been proposed in recent years that address the problem of model explainability and transparency. We present these novel

explanation approaches, characterize their strengths and limitations, position existing work with respect to the database (DB) community, and enumerate opportunities for data management research in the context of XAI.

### Representative Publications

Romila Pradhan, Aditya Lahiri, Sainyam Galhotra, Babak Salimi. Explainable AI: Foundations, Applications, Opportunities for Data Management Research. In Proceedings of the 2022 ACM International Conference on Management of Data (SIGMOD). 2022

Romila Pradhan, Aditya Lahiri, Sainyam Galhotra, Babak Salimi. Explainable AI: Foundations, Applications, Opportunities for Data Management Research. In Proceedings of the 38th IEEE International Conference on Data Engineering (ICDE). 2022

## Generating Interpretable Data-Based Explanations for Fairness Debugging

### PI: Romila Pradhan

### Current Students: Tanmay Surve

### Overview

A wide variety of fairness metrics and eXplainable Artificial Intelligence (XAI) approaches have been proposed in the literature to identify bias in machine learning models that are used in critical real-life contexts. However, merely reporting on a model's bias or generating explanations using existing XAI techniques is insufficient to locate and eventually mitigate sources of bias. We introduce Gopher, a system that produces compact, interpretable, and causal explanations for bias or unexpected model behavior by identifying coherent subsets of the training data that are root-causes for this behavior. Specifically, we introduce the concept of causal responsibility that quantifies the extent to which intervening on training data by removing or updating subsets of it can resolve the bias. Building on this concept, we develop an efficient approach for generating the top-k patterns that explain model bias by utilizing techniques from the machine learning (ML) community to approximate causal responsibility, and using pruning rules to manage the large search space for patterns.

### Representative Publications

Romila Pradhan, Jiongli Zhu, Boris Glavic, Babak Salimi. Interpretable Data-Based Explanations for Fairness Debugging. In Proceedings of the 2022 International Conference on Management of Data (SIGMOD), 2022.

Jiongli Zhu,Romila Pradhan, Boris Glavic, Babak Salimi. Demonstration of Generating Interpretable Data-Based Explanations for Fairness Debugging using Gopher. In Proceedings of the 2022 International Conference on Management of Data (SIGMOD), 2022.

Project URL: https://gopher-sys.github.io/

# Generating Interpretable Diagnostic Explanations for Black-box AI Systems

**PIs: Romila Pradhan, Babak Salimi (University of California San Diego), Boris Glavic (Illinois Institute of Technology, Chicago)**

## Overview

Artificial intelligence (AI) systems are increasingly deployed for decision-making in critical domains, such as healthcare, criminal justice, and finance. There is, however, growing concern that the opacity of these systems can perpetuate systemic biases and discrimination reflected in training data. Following increasing regulations by governmental agencies to generate human-understandable explanations for the behavior of these systems, the field of eXplainable Artificial Intelligence (XAI) witnessed a recent resurgence of interest. XAI tools are guided by social and ethical goals to: (a) increase societal acceptance of AI-based decision-making algorithms by establishing trust in outcomes, (b) provide users with actionable insights to change the results of algorithms in the future, and (c) enable the identification and debugging of sources of bias such as data collection strategies and training data that result in adverse and unexpected behavior. Existing approaches in XAI primarily focus on generating feature-based explanations that quantify the extent to which input feature values are responsible for the predictions of a model. These explanations suffer from severe limitations including the inability to capture causal relationships between variables and generating interventions that are not actionable in the real world. Furthermore, feature-based explanations are insufficient for generating diagnostic analyses that let practitioners trace unexpected or discriminatory algorithmic behavior back to training data. The algorithmic decisions could result from data errors and biases introduced during different stages of the analysis pipeline, such as data collection and preparation.

This project aims to address the limitations of existing approaches in XAI by developing novel techniques based on concepts from data management and machine learning to generate explanations for AI-based decision-making algorithms. We propose a two-pronged approach to reconcile the aforementioned objectives of XAI: first, we develop a novel framework based on probabilistic contrastive counterfactuals to provide insights into what causes the decisions of a black-box AI system, and second, we generate interpretable, diagnostic explanations for unexpected or discriminatory decisions made by the black-box AI system.

## Representative Publications

Sainyam Galhotra*, Romila Pradhan*, Babak Salimi. Explaining Black-Box Algorithms Using Probabilistic Contrastive Counterfactuals. In Proceedings of the 2021 International Conference on Management of Data (SIGMOD), 2021.

Project URL: https://lewis-system.github.io/

# Artificial Infrastructures

**PIs: Michael Salvo, John Sherrill, independent scholar**

**Current Students: Ean Hunt, CLA Wilke, Jefferey Chen, CLA Wilke, Eva Braumbauer, CLA Wilke, Meagan Hipsky, CLA Wilke**

## Overview

*Artificial Infrastructures* articulates the emergent roles of artificial intelligence specifically in high technology environments, with an eye towards the concerns of technical and professional writing experts, accessible to professionals with a wide range of proficiency. The book addresses the questions of What now? In an age when artificial agents draft text and respond to requests for unique prose. And what of the Turing Test?

*Artificial Infrastructures* develops a lasting argument about the nature of technology, ensuring it lasts longer than the current generation of AI tools. With a new GPT-3 (GPT-4) engine promised as this proposal is submitted, the next generation already promises to "radically disrupt" writing, and by extension, the lives and careers of writing professionals (and writing instructors & instruction). Rather than worry about the disruption of writing with the aid of technology, Artificial Infrastructures recognizes the "always already" nature of literacy and its technological enframement, arranged and presented for professional writers and professionals who write.

## Representative Publications

Sherrill, John T. & Salvo, Michael J. (2022). Automated Infrastructures: Participation's Changing Role in Postindustrial Work. Communication Design Quarterly (CDQ), vol. 10, no. 2, Sept. 2022, 22–31, https://doi.org/10.1145/3507857.3507860.

Sherrill, John T. & Salvo, Michael J. (2022). Distant Collaborations: Designing for Australia, Ireland, Qatar, and the USA. 2022 IEEE International Professional Communication Conference ProComm) Proceedings, IEEE Press, 2022, 154–59. ACM Digital Library, https://doi.org/10.1109/ProComm53155.2022.00031.

Michael J Salvo and John T Sherrill. 2024. A Research Ensemble of Humans, Machines, and Algorithms: Future Designs of Research and Scholarly Communication. In The 42nd ACM International Conference on Design of Communication (SIGDOC '24), October 20–22, 2024, Fairfax, VA, USA. ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3641237.3691666

# Testing Deep Learning Software

**PI: Lin Tan**

### Overview

Deep learning (DL) systems are widely used in domains including aircraft collision avoidance systems, Alzheimer's disease diagnosis, and autonomous driving cars. Despite the requirement for high reliability, DL systems are difficult to test.

Existing DL testing work focuses on testing the DL models, not the implementations (e.g., DL software libraries) of the models. One key challenge of testing DL libraries is the difficulty of knowing the expected output of DL libraries given an input instance. Fortunately, there are multiple implementations of the same DL algorithms in different DL libraries.

Thus, we propose CRADLE, a new approach that focuses on finding and localizing bugs in DL software libraries. CRADLE (1) performs cross-implementation inconsistency checking to detect bugs in DL libraries, and (2) leverages anomaly propagation tracking and analysis to localize faulty functions in DL libraries that cause the bugs. We evaluate CRADLE on three libraries (TensorFlow, CNTK, and Theano), 11 datasets (including ImageNet, MNIST, and KGS Go game), and 30 pre-trained models. CRADLE detects 12 bugs and 104 unique inconsistencies, and highlights functions relevant to the causes of inconsistencies for all 104 unique inconsistencies.

### Representative Publications

CRADLE: Cross-Backend Validation to Detect and Localize Bugs in Deep Learning Libraries. Hung Viet Pham, Thibaud Lutellier, Weizhen Qi, and Lin Tan. In the proceedings of the International Conference on Software Engineering. Acceptance Rate: 21% (109/529)

# RNCP: A Resilient Networking and Computing Paradigm for NASA Space Exploration

PI: Jin Wei-Kocsis

### Overview

There are three fundamental research objectives in this project: (1) to design a secure and decentralized computing and data analysis infrastructure, (2) to develop a data-driven resilient and cognitive networking management architecture, (3) to simulate, test, and validate our Resilient Networking and Computing Paradigm (RNCP). Our research with these three objects aims to advance the automation, environment-awareness, and intelligence of NASA's deep space network and to improve the resilience and scalability of the NASA space communication system. Furthermore, there is a possibility to integrate the work in the Cognitive Communications Project CubeSat Space mission to demonstrate the advanced cognitive networking architecture.

Project URL: https://polytechnic.purdue.edu/facilities/cyber-physical-social-systems-design-lab/research

# Generating Electricity Managed by Intelligent Nuclear Assets (GEMINA)

**PIs: Lefteri Tsoukalas, Alexander Heifetz**

**Current Students: Styliani Pantopoulou, Konstantinos Prantikos, Maria Pantopoulou**

*Funding Source: U.S. Department of Energy, Advanced Research Projects Agency-Energy (ARPA-E)*

## Overview

GEMINA aims to develop digital twin technology for advanced nuclear reactors and transform operations and maintenance (O&M) systems in the next generation of nuclear power plants. There is a need for tools that introduce greater flexibility in reactor systems, increased autonomy in operations, faster design iteration, and improved economic competitiveness. To accomplish this, we explore the application of data-driven and physics-informed models for monitoring sensor measurements and reactor transients. We also develop transfer learning methods which allow for monitoring with limited historical data, using models trained on different operating conditions.

## Representative Publications

Pantopoulou, S.; Ankel, V.; Weathered, M.T.; Lisowski, D.D.; Cilliers, A.; Tsoukalas, L.H.; Heifetz, A. Monitoring of Temperature Measurements for Different Flow Regimes in Water and Galinstan with Long Short-Term Memory Networks and Transfer Learning of Sensors. Computation 2022, 10, 108.

S. Pantopoulou, M. Weathered, D. Lisowski, L. H. Tsoukalas and A. Heifetz, "Temporal Forecasting of Distributed Temperature Sensing in a Thermal Hydraulic System With Machine Learning and Statistical Models," in IEEE Access, vol. 13, pp. 10252-10264, 2025.

Prantikos, K., Tsoukalas, L. H., & Heifetz, A. (2022). Physics-informed neural network solution of point kinetics equations for a nuclear reactor digital twin. Energies, 15(20), 7697.

Prantikos, K., Chatzidakis, S., Tsoukalas, L. H., & Heifetz, A. (2023). Physics-informed neural network with transfer learning (TL-PINN) based on domain similarity measure for prediction of nuclear reactor transients. Scientific reports, 13(1), 16840.

# EAGER: SaTC-EDU: Cybersecurity Education in the Age of Artificial Intelligence: A Novel Proactive and Collaborative Learning Paradigm

**PIs: Baijian Yang, Dr. Jin Wei-Kocsis (PI), Dr. Tonglin zhang (Co-PI), Dr. Baijian Yang (Co-PI)**

## Overview

This project is lead by Dr. Jin Kocsis

Artificial intelligence (AI) techniques, especially machine learning (ML), show great promise for improving quality of life. However, recent research has demonstrated that AI techniques can be manipulated, evaded, and misled. While progress has been made to better understand the trustworthiness and security of AI techniques, little has been done to translate this knowledge to education and training. There is a critical need to foster a qualified cybersecurity workforce that understands the usefulness, limitations, and best practices of AI technologies in the cybersecurity domain. This project will address this important issue by designing and implementing a virtual, proactive, and collaborative learning paradigm that can engage learners with different backgrounds. The approach will benefit a wide range of learners, especially underrepresented students. It will also help the general public understand the security implications of AI. This project has the ability to transform education at the intersection of cybersecurity and AI/ML; shed light on explainable AI in cybersecurity; and grow a cybersecurity workforce that possesses AI competencies. Products, including the research findings and curriculum, will be disseminated through a variety of mechanisms, such as workshops, peer-reviewed conferences, and journals.

This project builds research and education capacity through the formation of a multidisciplinary team with expertise in cybersecurity, AI, and statistics. The team will systematically investigate two cohesive research and education goals. First, an immersive learning environment will be developed to motivate students to explore AI/ML development in the context of real-world cybersecurity scenarios by constructing learning models with tangible objects. The proposed learning environment enables an AI/ML mechanism that will provide personalized explanations on the AI/ML outputs by considering the distinct background knowledge of the individual learners. Second, the team will design a proactive education paradigm encourages students to collaboratively identify new AI/ML-specific threats in the cybersecurity domain and develop innovative and trustworthy AI/ML solutions. The learning paradigm will ultimately enable effective retention and transfer of multidisciplinary AI-cybersecurity knowledge.

# Probabilistic Red-Teaming for Large Language and Vision-Language Models

**PI: Ruqi Zhang**

## Overview

As large language models (LLMs) and vision-language models (VLMs) grow more capable and widely deployed, they have also become increasingly susceptible to jailbreaks that bypass safety guardrails. Traditional red-teaming approaches often depend on heuristic search, genetic algorithms, or manually curated prompt pools, leading to limited coverage and poor scalability. These methods optimize adversarial examples one at a time, failing to capture the broader *distribution* of vulnerabilities that govern model behavior.

This project develops **probabilistic red-teaming**, a new framework that reframes adversarial prompt discovery as a problem of *probabilistic inference*. We introduce **VERA** (Variational infErence fRamework for jAilbreaking), which models jailbreak prompting as variational inference over the posterior distribution of adversarial prompts. A lightweight *attacker model* is trained to approximate this posterior, enabling it to efficiently generate diverse, high-quality jailbreak prompts for unseen queries without repeated optimization. By treating red-teaming as inference, VERA captures the underlying uncertainty and structure of vulnerabilities, providing a richer characterization of model failure modes.

Extending this approach to multimodal models, **VERA-V** formulates jailbreak discovery as learning a joint distribution over coupled text-image prompts. This probabilistic perspective allows for coordinated attacks that combine linguistic and visual perturbations to evade detection and safety filters. VERA-V integrates complementary mechanisms, typography-based text embedding, diffusion-guided adversarial image synthesis, and structured visual distractors, to fragment model attention and expose hidden weaknesses in visual reasoning. Empirical results demonstrate substantial improvements over existing methods.

Together, these efforts establish a principled, scalable foundation for red-teaming large language and multimodal models. By combining probabilistic inference with adversarial testing, this project aims to (1) advance the science of model evaluation under uncertainty, (2) provide actionable insights into real-world vulnerabilities, and (3) lay the groundwork for *trustworthy, probabilistically aligned AI systems*.

## Representative Publications

VERA: Variational Inference Framework for Jailbreaking Large Language Models. Neural Information Processing Systems (NeurIPS), 2025

VERA-V: Variational Inference Framework for Jailbreaking Vision-Language Models. Preprint, 2025

# Inter-Play Between Cyberscurity and Deep-Learning

PI: Xukai Zou, Dr. Feng Li

**Current Students: Agnideven Sundar, Ryan Hosler**

**Overview**

This work focuses the combination and inter-play of cybersecurity and artificial intelligence. On one hand, the project investigates security issues by utilizing deep-learning techniques such as Unsupervised Deep Learning for Android Malware Feature Extraction and Detection Using BiGAN. On the other hand, the project investigates security and privacy issues associated with deep-learning models, such as Backdoor Attack, Detection, and Defense in Decentralized Federated Learning for Non-IID Data.

# Assured Identity and Privacy

## Directed Infusion of Data

**PI: Hany Abdel-Khalik**

**Current Students: Arvind Sundaram, Tyler Lewis, Chloe Yoder**

Funding Source: Idaho National Laboratory

Overview

The Directed Infusion of Data (DIOD) paradigm is a novel data-based obfuscation procedure developed in response to growing data privacy concerns in wake of the rise in complexity, scale, and capability of artificial intelligence and machine learning (AI/ML) tools. General data sharing and collaboration typically requires proprietary data transfer, i.e., a stakeholder hands their data, usually in an encrypted form, to a data analyst; though all parties are generally considered trustworthy, data privacy is put at unnecessary risk simply by its distribution, thereby endangering financial resources, personal security, and proprietary/ classified material. The key issue with ensuring data privacy is that the data need to be protected while retaining their utility; many proposed methods enforce limiting conditions that avoid sharing the data but sacrifice some of its utility in doing so, thus limiting the capability of the analyst during collaboration.

The DIOD paradigm seeks to obfuscate data in an efficient manner while allowing for both data security as well as utility. By obfuscating the dynamic behavior of proprietary data with that of an unrelated dataset, the inference provided by the true data, e.g., classification, presence of anomalies, or variable dependencies, may be preserved in the new, obfuscated set of data. Using DIOD's form of obfuscation, the data remain usable for the desired purpose, but the dynamic behavior, i.e., the 'identity', is changed so that proprietary details cannot be reverse engineered, thus mitigating the need to risk vital information during collaboration or outsourced computation.

A secondary benefit of DIOD is the flexibility with which the data can be masked; in order to add an additional layer of security, the structure of the data can be altered, i.e., timeseries data obfuscated as image-based data, thereby allowing data masking well-suited to the needs of a particular analysis.
Representative Publications

Arvind Sundaram, Hany Abdel-Khalik & Ahmad Al Rashdan, "Deceptive Infusion of Data: A Novel Data Masking Paradigm for High-Valued Systems", Nuclear Science and Engineering, Feburary 2022

## Obfuscation of Audio Signals

**PI: Hany Abdel-Khalik**

**Current Students: Chloe Yoder, Tyler Lewis, Arvind Sundaram**

**Overview**
The current industrial paradigm has seen an increased variety in data volume, real-time analytics, and

the role of artificial intelligence and machine learning (AI/ML) tools that process many types of data, including timeseries, image, text-based, and audio signals to inform operational decisions. While the corresponding rise in data utility has allowed for efficient, safe, and cheap data analytics, these tools often invite major security concerns in light of the increased risk of intrusion detection and data leakage. Among the chief security challenges is that of audio-based signals; whether this be machine-operation signals, proprietary conversations, etc., audio data may overtly detail sensitive information that necessitates a specialized form of obfuscation to ensure are protected in real-time.

In the case of data sharing, security is typically ensured by encryption-based techniques which require the genuine audio signal to be released. The so-called issue of trust can be solved by obfuscating the real signal with an arbitrary signal or image such that its identity is no longer at risk, thereby mitigating the uncertainties associated with typical collaboration. The Directed Infusion of Data (DIOD) paradigm can be directly applied to audio-based data, thus allowing a typical conversation to take the form of a song, another conversation, or an arbitrary signal.

A secondary benefit of the DIOD paradigm applied to audio-based data is that direct information of the Fourier-type analyses, i.e., the most fundamental frequencies of the audio data, or time-stamped signal properties, may be preserved to suit the needs of a downstream analysis, thus allowing covert data masking of full conversations while retaining characteristic attributes.

## Foundations of Cyber-Physical Infrastructure for Creative Design and Making of Cyber-physical Products

**PIs: Mike Atallah, PI is Jitesh Panchal, co-PIs are Mikhail Atallah and Karthik Ramani**

### Overview

The focus in this proposal is on the emerging paradigm of democratization of the product innovation process, especially by engaging a vastly larger pool of talent to generate new ideas which in turn will create new cyber-physical products. The primary objective is to lay the foundations of a cyber-physical infrastructure for the creative design and making of realizable products. This will be done by addressing fundamental barriers to participation, model-based engineering, and information sharing. The information-security facet of the project aims to overcome information-related impediments to collaboration and information sharing, and thereby enable new forms of interactions between individuals, the physical world, and the computational aspects of product realization.

### Representative Publications

Shumiao Wang, Siddharth Bhandari, Mikhail Atallah, Jitesh Panchal, and Karthik Ramani. Secure Collaborations in Engineering System Design, Proc. of ASME 2014 International Design and Engineering Technical Conference and Computers and Information in Engineering Conference (IDETC/CIE 2014), Buffalo, New York, August 2014. [This is the flagship conference for design engineering.]

# Customized Privacy Mechanisms for Statistical Inference

**PI: Jordan Awan, Guang Cheng, Professor of Statistics and Data Science, University of California, Los Angeles, Salil Vadhan, Vicky Joseph Professor of Computer Science and Applied Mathematics, Harvard University, Aleksandra Slavkovic, Professor of Statistics, Penn State**

**Current Students: Zhanyu Wang (graduated), Yuki Ohnishi, Yue Wang**

## Overview

Differential privacy (DP) is the state-of-the-art framework for formal privacy protection, but many available DP methods are designed primarily for estimation. On the other hand, in many scientific problems, it is important to have a complete statistical analysis, which may include 1) a particular statistical model, 2) estimation, and 3) uncertainty quantification (such as confidence intervals and hypothesis tests). In this project, we design DP mechanisms specifically for these statitistical tasks, focusing primarily on the uncertainty quantification. One general technique we explore is the use of the bootstrap in combination with a privacy mechanism to understand the sampling distribution of the private summaries. Besides general statistical applications, we also study the particular problem of valid causal inference from both randomized and observational studies.

## Representative Publications

Ohnishi, Yuki, and Jordan Awan. "Locally Private Causal Inference for Randomized Experiments." *arXiv preprint arXiv*:2301.01616 (2023).

Awan, Jordan, and Salil Vadhan. "Canonical noise distributions and private hypothesis tests." *The Annals of Statistics* 51, no. 2 (2023): 547-572.

Wang, Zhanyu, Guang Cheng, and Jordan Awan. "Differentially private bootstrap: New privacy analysis and inference strategies." *arXiv preprint arXiv*:2210.06140 (2022).

Awan, Jordan, and Yue Wang. "Differentially Private Kolmogorov-Smirnov-Type Tests." *arXiv preprint arXiv*:2208.06236 (2022).

Awan, Jordan, and Aleksandra Slavković. "Differentially private inference for binomial data." *Journal of Privacy and Confidentiality* 10, no. 1 (2020): 1-40.

Awan, Jordan, and Aleksandra Slavković. "Differentially private uniformly most powerful tests for binomial data." *Advances in Neural Information Processing Systems* 31 (2018).

# Differential Privacy Methods for Machine Learning and Complex Data Structures

**PI: Jordan Awan, Matthew Reimherr, Principal Research Scientist at Amazon and an Affiliate Professor of Statistics at Penn State, Aleksandra Slavkovic, Professor of Statistics, Penn State, Vinayak Rao, Associate Professor of Statistics, Purdue University**

**Current Students: Taegyu Kang, Sehwan Kim, Jinwon Sohn, Ana Kenney**

### Overview

As more personal data is collected and analyzed, there is a growing need for formal privacy protection. Differential privacy (DP) has arisen as the state-of-the-art method in privacy protection, but many DP methods are limited to simplistic settings and are not optimized for complex machine learning tasks. In this project, we develop and optimize DP algorithms for various machine learning tasks which can analyze complex datasets. Specifically, we develop DP methods for 1) empirical risk minimization (which encompases a wide variety of machine learning methods), 2) functional data analysis, and 3) topological data analysis.

### Representative Publications

Kang, Taegyu, Sehwan Kim, Jinwon Sohn, and Jordan Awan. "Differentially Private Topological Data Analysis." *arXiv preprint arXiv*:2305.03609 (2023).

Awan, Jordan, and Vinayak Rao. "Privacy-aware rejection sampling." *Journal of machine learning research* 24, no. 74 (2023): 1-32.

Awan, Jordan, Ana Kenney, Matthew Reimherr, and Aleksandra Slavković. "Benefits and pitfalls of the exponential mechanism with applications to hilbert spaces and functional pca." In *International Conference on Machine Learning*, pp. 374-384. PMLR, 2019.

Reimherr, Matthew, and Jordan Awan. "KNG: The K-norm gradient mechanism." *Advances in neural information processing systems* 32 (2019).

Reimherr, Matthew, and Jordan Awan. "Elliptical perturbations for differential privacy." *Advances in Neural Information Processing Systems* 32 (2019).

# Foundations of Differential Privacy

**PI: Jordan Awan, Aleksandra Slavkovic, Professor of Statistics, Penn State, Salil Vadhan, Vicky Joseph Professor of Computer Science and Applied Mathematics, Harvard University**

**Current Students: Aishwarya Ramasethu, Young Hyun Cho**

### Overview

Differential privacy (DP) has arisen as the state-of-the-art framework for formal privacy protection when analyzing sensitive data. However, the fundamentals of DP are still not well understood. There are many

important questions for DP, which with proper understanding can help researchers understand the fundamental limits of what can be achieved with privacy protection as well as the optimal methods to achieve DP. In this project, we study the foundations of DP such as 1) the choice of privacy definition, 2) optimal noise to achieve a given DP definition, and 3) understanding the sensitivity of a statistic and how this relates to how much noise is required to protect privacy.

### Representative Publications

Awan, Jordan, and Aishwarya Ramasethu. "Optimizing Noise for $ f $-Differential Privacy via Anti-Concentration and Stochastic Dominance." *arXiv preprint arXiv*:2308.08343 (2023).

Awan, Jordan, and Salil Vadhan. "Canonical noise distributions and private hypothesis tests." *The Annals of Statistics* 51, no. 2 (2023): 547-572.

Awan, Jordan, and Jinshuo Dong. "Log-concave and multivariate canonical noise distributions for differential privacy." *Advances in Neural Information Processing Systems* 35 (2022): 34229-34240.

Awan, Jordan, and Aleksandra Slavković. "Structure and sensitivity in differential privacy: Comparing k-norm mechanisms." *Journal of the American Statistical Association* 116, no. 534 (2021): 935-954.

## Statistical Methods for Privatized Data

**PIs: Jordan Awan, Robert Molinari, Assistant Professor, Department of Statistics, Auburn University, Nianqiao Ju, Assistant Professor, Department of Statistics, Purdue University, Vinayak Rao, Associate Professor, Department of Statistics, Purdue University, Ruobin Gong, Assistant Professor, Department of Statistics, Rutgers University, Andres Felipe Barrientos, Assistant Professor of Statistics, Florida State University**

**Current Students: Zhanyu Wang (now graduated), Yu Wei Chen, Xinlong Du, Aidan Davis, Samuel Forfang**

### Overview

As more personal data is collected, analyzed, and published by tech companies, academic researchers, and government agencies, the concern for privacy protection increases. To address these concerns, formal privacy protection methods, such as differential privacy (DP), are becoming widely employed by tech companies as well as federal statistical agencies. To protect privacy, DP methods require the introduction of additional randomness into the analyses, which “covers up” what any individual has contributed to the database. However, this extra noise results in additional bias and variance. Due to this, researchers and policy makers who depend on these products have found that traditional statistical tools give misleading and unreliable results. For example, the 2020 US Decennial Census employed differential privacy and social scientists, demographers, and economists have raised alarms over the difficulty in obtaining valid inferences.

This raises the question: how can we enable researchers to obtain valid statistical inference on privatized data for a variety of models and privacy mechanisms? While previous work has begun developing statistical tools in DP, most of the prior solutions are tailored to narrow problems, and often lack concrete statistical guarantees. We propose to develop general-purpose statistical inference methods, which are

applicable to a wide variety of analyses on privatized data, and which are based on the leading methods of likelihood-free inference. We will also develop new privacy-protecting procedures that will be optimized to enable reliable statistical inferences, including unbiased estimators, valid confidence intervals, calibrated hypothesis tests, and posterior inference.

**Intellectual Merit**
The proposed interdisciplinary research will deliver both theoretical and practical tools for the advancement of statistical approaches in complex settings such as those entailed by the added noise of DP mechanisms. The likelihood-free methods that will be studied and implemented from the proposed research will provide computationally efficient solutions to solve complex problems in both the private and in the non-private data settings. These will be among the first general tools to perform adequate statistical inference for DP outputs that are not tailored to very specific statistical tests with narrow applications, and can be broadly applied to many statistical procedures that are commonly used in the social sciences and other fields of research. In addition, studying these approaches will open future avenues of research to efficiently obtain statistical outputs with adequate finite sample and asymptotic properties (unbiasedness, consistency, efficiency) in settings for which solutions do not currently exist, such as reliable inference procedures for the common cases of missing, censored or truncated data. More specifically, the study of how certain statistics can deliver better results within a simulation-based framework can greatly contribute to the development of techniques such as co-sufficient sampling and indirect inference whose theoretical and practical advantages have not been fully exploited.

### Representative Publications

Awan, Jordan, and Zhanyu Wang. "Simulation-based, finite-sample inference for privatized data." *arXiv preprint arXiv*:2303.05328 (2023).

Ju, Nianqiao, Jordan Awan, Ruobin Gong, and Vinayak Rao. "Data augmentation MCMC for bayesian inference from privatized data." *Advances in neural information processing systems* 35 (2022): 12732-12743.

Awan, Jordan, Andres Felipe Barrientos, and Nianqiao Ju. "Statistical Inference for Privatized Data with Unknown Sample Size" *arXiv preprint arXiv:2406.06231* (2024).

## ABAUS: Active Bundle AUthentication Solution Based on SDN for Vehicular Networks

### PI: Bharat Bhargava

### Current Students: Aala Oqab Alsalem

### Overview
Vehicular ad hoc networks (VANETs) are gaining more and more interest in intelligence transportation system research fields. They allow optimized traffic management due to improved vehicle resource usage and real-time information exchanges. However, being in an open environment introduces different security and privacy challenges. Attackers can sniff radio signals and forge the transmitted information

leading to sensitive data leaking or compromising. This paper examines the preservation of privacy information in VANET communications. We use the Active Bundle (AB) for vehicle authentication and data preservation based on software-defined networks (SDNs). Our proposal benefits from the SDN infrastructure to guarantee fluent centralized management while using the AB guarantees data integrity and confidentiality throughout its entire lifecycle. Analytical studies and simulations show that our solution efficiently preserves VANET users' privacy with minimal effects on network transmission quality.

Project URL: https://ieeexplore.ieee.org/document/10460563

## An Attack to One-Tap Authentication Services in Cellular Networks

**PIs: Bharat Bhargava, Z. Cui, B. Cui, J. Fu**

### Overview

The One-Tap Authentication (OTAuth) based on the cellular network is a password-less login service provided by Mobile Network Operator (MNO) through the unique communication gateway access technique. The service allows app users to quickly sign up or log in with their mobile phone numbers without entering a password. Due to its convenience, OTAuth has been widely used by various apps. However, some studies have elaborated that OTAuth services are of great drawbacks from the perspective of mobile security and identified several flawed designs, which make the MNO cannot distinguish malicious apps from normal ones and cause impersonation attacks. In this paper, we further analyze OTAuth services from the perspective of 4G and 5G cellular networks and focus on two important procedures in which the cellular network plays an important role in OTAuth services. Not surprisingly, we discover a new fundamental design flaw in determining whether the runtime environment supports OTAuth services. Moreover, we propose a mature attack paradigm by exploiting this flaw, which allows an attacker to login or register one app as a victim. To evaluate the impact of the attack, we have examined 100/90/100 Android/iOS/HarmonyOS apps for OTAuth services of 3 mainstream MNOs in China. The experimental results show that our proposed attack is applicable to almost all the apps that support OTAuth services, and affects more apps than the attacks that have been reported before. Finally, we propose several countermeasures to defend against the attack. Note that, for security's sake, we have already reported our findings to authorized parties and received their confirmations.

### Representative Publications

Z. Cui, B. Cui, J. Fu and B. K. Bhargava, "An Attack to One-Tap Authentication Services in Cellular Networks", IEEE Transactions on Information Forensics and Security, vol. 18, pp. 5082-5095, 2023, doi: 10.1109/TIFS.2023.3304840

Project URL: https://www.cs.purdue.edu/homes/bb/zhiwei1.pdf

# Privacy – Preserving Data Dissemination in Untrusted Cloud

**PI: Bharat Bhargava**

**Current Students: Aala Oqab Alsalem**

**Overview**
B2B (business-to-business) systems often use service-oriented architecture (SOA) with decomposed business services. These services can interact and share data among each other. Service might use a cloud – hosted database, such as a non - relational encrypted key – value store. However, the cloud platform hosting the database can be untrusted. Data owner needs to be sure that each service can access only those segments of a shared database for which the service is authorized. Furthermore, data requests can come from a service also hosted by untrusted cloud. Hence, there is a need for designing a cloud enterprise framework that can ensure privacy-preserving data dissemination in SOA and accurately detect data leakages. We design and prototype a solution that ensures privacy – preserving dissemination of data. The solution is based on (a) role-based access control, (b) cryptographic capabilities of client's browser, (c) authentication method, (d) subject's trust level. The prototype enables privacy – preserving dissemination of Electronic Health Records (EHRs) hosted in an untrusted cloud. Keywords—privacy; trust management; data dissemination; access control; SOA; database privacy; cloud computing

**Representative Publications**
IEEE Cloud Computing Conference

Project URL: https://www.cs.purdue.edu/homes/bb/#research Secure

# QPCASIN: A Quantum-Defended Privacy-Aware Preemptive Handover-Enabled Continuous Authentication in Space Information Networks

**PIs: Bharat Bhargava, B. Palaniswamy, A. Karati, T. -Y. Chen, A. K. Das**

**Overview**
The Space Information Network (SIN) plays a crucial role in terrestrial communication, delivering time-bound services from ground stations to users. It relies on moving low-orbit earth (LEO) satellites for uninterrupted coverage. However, untrustworthy connectivity poses several security challenges during handover services for users maintained by the satellites. While traditional cryptographic techniques provide a degree of security, the advent of quantum computing exposes significant vulnerabilities. This work proposes a quantum-safe and continuous authentication mechanism with handover provision. The proposed authentication protocol uses post-quantum primitives of the Frodo key encapsulation mechanism, currently an approved mechanism under ISO/IEC 18033-2. It ensures privacy and ensures users' anonymity. The security of the proposed protocol is analyzed using the quantum random oracle (QROM) model. Formal verification confirms its safety for practical adoption as a post-quantum candidate. Further, the performance evaluation shows an authentication delay and energy consumption of

the proposed protocol within practical limits, making it a suitable candidate for privacy-preserved post-quantum adoption for SIN.

### Representative Publications

B. Palaniswamy, A. Karati, T. -Y. Chen, A. K. Das and B. K. Bhargava, "QPCASIN: A Quantum-Defended Privacy-Aware Preemptive Handover-Enabled Continuous Authentication in Space Information Networks," in *IEEE Transactions on Information Forensics and Security*, vol. 20, pp. 6910-6922, 2025, doi: 10.1109/TIFS.2025.3583246. keywords: {Authentication;Protocols;Satellites;Security;Handover;Low earth orbit satellites;Cryptography;Quantum computing;Privacy;Space vehicles;Space information networks;post-quantum security;lattice cryptography;authentication;handover mechanism;formal verification;anonymity},

Project URL: https://ieeexplore.ieee.org/document/11050854

## Formal Privacy for Multi-stage Probability Samples

**PIs: Chris Clifton, Keith Merrill (Brandeis University)**

**Current Students: Shawn Merrill, Eric Hanson (Brandeis University)**

### Overview

*Funding Source: U.S. Census Bureau*

Census products based on sampling, such as the U.S. Census Burueau's American Community Survey, use a complex weighting process to improve accuracy. We have been studying how to apply formal privacy methods to weighted samples. A straightforward application of differential privacy results in unacceptably high noise: as the potential exists for a highly weighted individual to have a huge impact on results, considerable noise is needed to protect against this possibility.

This project is developing techniques to address these issues: providing formal privacy guarantees while releasing data with quantifiable data value.

## Securing the Software Supply Chain: Theories, Measurements, Runtime Defenses, and Software Signing Infrastructure (PKI) for commercial and open-source software

**PIs: Jamie Davis, Santiago Torres-Arias**

*Funding Source: Cisco, National Science Foundation (NSF)*

### Overview

Many software applications incorporate third-party packages distributed by package registries. Guaranteeing package provenance – knowledge of authorship – along this supply chain is a necessary part of ensuring that the software applications that run our societies are trustworthy. Although

package maintainers can guarantee package authorship through software signing based on public-key cryptography, the adoption of signing has been slow. Many prior works have discussed challenges with the different generations of signing tools. However, recent cyberattacks have prompted a renewed emphasis on software signing from technology leaders such as Google (SLSA) and the US NIST (NIST SP 800-204D).

One perennial problem with the adoption of software signing has been the myriad competing signing tools and inconsistency across different package registries. The goal of this project is to provide a theoretical and empirical basis to understand **what factors limit and predict the adoption of software signing in open-source software in order to promote the adoption of this practice**. We are applying both quantitative methods (mining software repositories for hypothesis testing) and qualitative methods (human factors -- interviews, surveys). Building on this, we are developing the Sigstore signing platform with a goal of substantially solving the signing problem.

### Representative Publications

*Signing in Four Public Software Package Registries: Quantity, Quality, and Influencing Factors*.
Schorlemmer, Kalu, Chigges, Ko, Ishgair, Bagchi, Torres-Arias, and **Davis**.
Proceedings of the 45th IEEE Symposium on Security and Privacy (IEEE S&P) 2024.

*SoK: Analysis of Software Supply Chain Security by Establishing Secure Design Properties*.
Okafor, Schorlemmer, Torres-Arias, and **Davis**. Proceedings of the 1st ACM Workshop on Software Supply Chain Offensive Research and Ecosystem Defenses (SCORED) 2022.

*An Empirical Study of Artifacts and Security Practices in the Pre-trained Model Supply Chain*.
Jiang, Synovic, Sethi, Indarapu, Hyatt, Schorlemmer, Thiruvathukal, and **Davis**. Proceedings of the 1st ACM Workshop on Software Supply Chain Offensive Research and Ecosystem Defenses (SCORED) 2022.

*ZTD-JAVA: Mitigating Software Supply Chain Vulnerabilities via Zero-Trust Dependencies*.
Amusuo, Robinson, Singla, Peng, Machiry, Torres-Arias, Simon, and **Davis**.
Proceedings of the ACM/IEEE 47th International Conference on Software Engineering (ICSE) 2025.

*An Industry Interview Study of Software Signing for Supply Chain Security*.
Kalu, Singla, Okafor, Torres-Arias, and **Davis**. arXiv 2024.

*Why Johnny Signs with Next-Generation Tools: A Usability Case Study of Sigstore*.
Kalu, Okorafor, Singla, Torres-Arias, and **Davis**

*Why Software Signing (Still) Matters: Trust Boundaries in the Software Supply Chain*.
Kalu and **Davis**

*DiVerify: Hardening Identity-Based Software Signing with Programmable Diverse-Context Scopes*.
Okafor, **Davis**, and Torres-Arias.

*ARMS: A Vision for Actor Reputation Metric Systems in the Open-Source Software Supply Chain*.
Kalu, Okorafor, Durak, Laine, Moreno, Torres-Arias, and **Davis**.

Project URL: https://www.sigstore.dev

# Language Support for Precise Privacy-Preserving Computation

**PIs: Benjamin Delaware, Milind Kulkarni**

**Current Students: Raghav Malik, Qianchuan Ye**

## Overview

Protecting the privacy of personal information is not just a matter of ethics, but of law: HIPAA and FERPA, for example, restrict how an individual's private healthcare and educational records can be used and released. These laws intersect in a delicate way with applications of modern data analyses that depend on protected data. In the healthcare setting, for example, sophisticated machine-learned models have the potential to revolutionize the diagnosis, prevention, and treatment of complex diseases like cancer. In this scenario, the data being analyzed is the very information that cannot be shared, including, potentially, with the developer of the proprietary diagnostic model.

The development of usable tools and techniques for expressing these sorts of privacy-preserving computations has been a topic of intense interest for several years, spurred in part by advances in cryptographic protocols and programming languages for secure multiparty computation (MPC). Despite recent progress, the current state of the art still suffers from some important limitations that hinder their adoption:

- Systems for privacy-preserving machine learning (e.g., Microsoft Research's CHET framework) have mainly targeted neural network–based approaches, with less consideration for alternative approaches like decision forests. A key challenge is how to support the wider variety of computational patterns that can arise in more general programs in an efficient manner.
- There is a considerable gap between the high-level semantic privacy requirements of the user and the guarantees provided by the underlying system. This forces the developer to ensure, for example, that the bits of information leaked by a computation are in compliance with the legal demands of HIPAA.
- There is often a fundamental tension between the efficiency of a particular computation and how much private information it leaks, and the programmer is responsible for exploring these tradeoffs. In the face of complex models that depend on different types of private data with different privacy requirements, providing efficiency while guaranteeing privacy is a daunting task.

This project aims to develop techniques for building efficient, general, privacy-preserving computations that tackle these key challenges. This high-level goal is divided into three complementary thrusts:

- A semantics-based approach to specifying privacy guarantees for secure computations (e.g., maybe it is safe to release a patient's lab results as long as no personally identifying information is also leaked, but if PII is leaked, other medical data cannot be), as well as language support for writing data-intensive privacy-preserving models that use these specifications. Leveraging our previous work on oblivious high-level programming languages, we automatically synthesize efficient data structures and programs for performing privacy-preserving machine learning.
- Automatic generation of mixed-mode computations that selectively release "safe" information to generate more efficient programs. We will build on our work for compiler optimization of FHE programs to create efficient, vectorized implementations of these programs.
- An investigation of how to automatically synthesize new programs that trade off information leakage for efficiency, while still meeting the necessary privacy requirements.

### Representative Publications

Yuyan Bao, Kirshanthan Sundararajah, Raghav Malik, Qianchuan Ye, Christopher Wagner, Nouraldin Jaber, Fei Wang, Mohammad Hassan Ameri, Donghang Lu, Alexander Seto, Benjamin Delaware, Roopsha Samanta, Aniket Kate, Christina Garman, Jeremiah Blocki, Pierre-David Letourneau, Benoit Meister, Jonathan Springer, Tiark Rompf, and Milind Kulkarni. 2021. HACCLE: metaprogramming for secure multi-party computation. In Proceedings of the 20th ACM SIGPLAN International Conference on Generative Programming: Concepts and Experiences (GPCE 2021). Association for Computing Machinery, New York, NY, USA, 130–143. https://doi.org/10.1145/3486609.3487205

Qianchuan Ye and Benjamin Delaware. 2022. Oblivious algebraic data types. Proc. ACM Program. Lang. 6, POPL, Article 51 (January 2022), 29 pages. https://doi.org/10.1145/3498713

Qianchuan Ye and Benjamin Delaware. 2023. Taype: A Policy-Agnostic Language for Oblivious Computation. Proc. ACM Program. Lang. 7, PLDI, Article 147 (June 2023), 25 pages. https://doi.org/10.1145/3591261

Raghav Malik, Vidush Singhal, Benjamin Gottfried, and Milind Kulkarni. 2021. Vectorized secure evaluation of decision forests. In Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation (PLDI 2021). Association for Computing Machinery, New York, NY, USA, 1049–1063. https://doi.org/10.1145/3453483.3454094

Raghav Malik, Kabir Sheth, and Milind Kulkarni. 2023. Coyote: A Compiler for Vectorizing Encrypted Arithmetic Circuits. In Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 3 (ASPLOS 2023). Association for Computing Machinery, New York, NY, USA, 118–133. https://doi.org/10.1145/3582016.3582057

## Standoff Inverse Analysis and Manipulation of Electronic Systems

**PIs: Edward Delp, Kevin G. Gard (North Carolina State University), Andreas Cangellaris (University of Illinois at Urbana-Champaign), Lawrence Carin (Duke University), William J. Chappell (Purdue University), Michael Steer (North Carolina State University)**

### Overview

*Funding Source: U.S. Army Research Laboratory (ARL)*

Given the wide use of RF devices for applications ranging from data networks to wireless sensors, it is of interest to identify the types of devices that are located in an environment. In order to locate and characterize RF devices, the environment must be probed. This becomes the problem of determining the properties of an RF circuit by sending it a carefully designed signal and examining the returned signal. The device is then authenticated by identifying certain characteristics of the return signal. Identification through “forensic characterization” means identifying the type of RF device, make, model, configuration, and other characteristics of the device based on observations of the data that the device produces.

The characterization framework is used to classify signals emitted from an RF device that are generated in response to a probe signal. Once transmitted into the environment, the probe signal is received by the antenna of the RF device and sent to several circuit components such as filters and amplifiers. Each circuit component modifies the probe signal, thereby embedding information into the signal. The embedded information is unique to each circuit component and when input into a nonlinear component such as an amplifier, a percentage of the probe signal is reflected. The reflected signal is re-radiated from the RF device and measured. Once measured, a forensic signature is extracted from the re-radiated signal and characterized by a classification system.

### Representative Publications

A. K. Mikkilineni, D. King-Smith, S.B. Gelfand, and E.J. Delp, "Forensic Characterization of RF devices," First IEEE International Workshop on Information Forensics and Security, WIFS 2009, pp.26-30, 6-9 Dec. 2009.

Project URL: http://www.siames.org

## Effects of Body Position on Facial Recognition in Police Body-Worn Cameras

**PI: Stephen Elliott, Dr. Kathryn Seigfried-Spellar**

**Current Students: Julia Bryan**

### Overview

The purpose of this project is to identify the effects of a police officer's body position on facial recognition identification accuracy in body-worn cameras (BWCs). More specifically, this research will examine how three different manipulations of a single policing stance will affect the facial recognition comparison scores for a known target using footage captured from a popular BWC under controlled conditions.

## Evaluation of Face Recognition for Law Enforcement: Book-In Photo Acquisition and Surveillance Camera Placement

**PIs: Stephen Elliott, Kevin O'Connor**

**Current Students: Holbrook Hankinson, Filippo Lippi, Deidra Wynn, Yuyuan Liu, Andres Rosas-Vargas**

### Overview

The purpose of our research is to determine the best way possible to provide quality facial recognition to local gas stations and local law authorities. Currently, many gas stations do not have their surveillance cameras in optimal positions for local law enforcement to accurately identify persons of interest Our goal was to assist in this process by conducting trials using our own faces and cameras. Within these trials, we were able to determine what positioning would be optimal for image quality, as well as testing the difference of light and distance that would also affect image quality scoring. The research was conducted

at a local Marathon in West Lafayette, IN.

## Post-Mortem Biometrics

**PIs: Stephen Elliott, Kevin O'Connor**

**Current Students: Joe Zweng**

**Overview**
The purpose of this research is to determine if commercially available biometric systems, specifically fingerprint, and iris, can be utilized on post-mortem individuals. In the event that the identity of a decedent is unknown, it can be difficult for law enforcement to correctly identify the subject. Biometrics could be used if images captured met or exceeded the image quality threshold and the necessary features could be extracted from the image. To complete this research, the International Center for Biometric Research has partnered with the Lafayette Police Department.

## A Comprehensive Approach for Data Quality and Provenance in Sensor Networks

**PIs: Sonia Fahmy, Elisa Bertino**

**Current Students: Salmin Sultana, S. M. Iftekharul Alam**

**Overview**
*Funding Source: National Science Foundation (NSF)*
Sensor networks are being increasingly deployed in a large variety of domains, ranging from critical IT infrastructures to computational medicine. Sensor networks allow real-time gathering of large amounts of data that can be mined and analyzed for taking critical actions. As such, sensor networks are a key component of any decision-making infrastructure. A critical issue in this context is the trustworthiness of the data being collected. Data integrity and quality decide the trustworthiness of data. Without integrity, any information extracted from the available data cannot be trusted. Data integrity can be undermined not only by errors introduced by users, measurement devices and applications, but also by malicious subjects who may inject inaccurate data with the goal of deceiving the users of the data. Therefore, it is critical that data integrity issues, including how to measure data quality and how to use data provenance for integrity assessment, be investigated for multi-sensor data integration, situation assessment, and numerous other functions. A fundamental tradeoff exists between data quality and the cost to gather and protect this data, e.g., in terms of sensor node energy.

The objective of the proposed research is to design and develop a comprehensive approach to the problem of assessing integrity of continuous data streams in sensor networks, taking into account cost and energy constraints.

**Representative Publications**
S. Sultana, M. Shehab, E. Bertino, "Provenance based Mechanism to Identify Malicious Packet Dropping Adversary in Sensor Network", In Proceedings of the 8th Workshop on Wireless Ad hoc and Sensor

Networks (WWASN 2011).

# Differentially Private Data Synthesis: Practical Algorithms and Statistical Foundations

**PI: Ninghui Li**

**Current Students: Yuntao Du**

*Funding Source: National Science Foundation (NSF)*

## Overview

Funded by National Science Foundation: **Collaborative Research: SaTC: CORE: Small: Differentially Private Data Synthesis: Practical Algorithms and Statistical Foundations**. 07/05/2023 - 06/30/2026.

Data collected by organizations and agencies are a key resource in today?s information age and fuel a significant part of today's economy. However, the disclosure of those data poses serious threats to individual privacy. One important approach to using data while protecting privacy is differential private data synthesis (DPDS). That is, given as input a private dataset, one uses a differentially private algorithm to generate synthetic datasets that are ?similar? to the input dataset. While DPDS has received much attention in recent years, our understanding on this topic remains limited. This project takes a multi-disciplinary approach to advance our scientific understanding as well as improve practice techniques for DPDS. More specifically, this project?s novelties are as follows. First, it systematically explores the design space in marginal-based DPDS algorithms that have been proven to be effective in NIST competitions on DPDS, while also taking insights from data synthesis techniques developed in similar fields (often not satisfying DP). Second, it develops statistical theories that both are motivated by the empirical performances of DPDS algorithms, and guide the empirical research of these algorithms.

The project?s broader significance and importance are as follows. We are in the information economy. Data of all kinds, such as online interaction, medical sensor data, genomic data, and location data are being collected. Practical techniques that enable use of these data while protecting individual privacy are crucially needed and will greatly enhance the value of such data. Users will gain from increased control of their private information, and society as a whole will benefit from deriving maximal benefit from aggregated data. PIs plan to jointly develop and teach a graduate-level course on synthetic data based on the existing research in this area as well as research results from this project, and involve undergraduate students in research.

This project has two thrusts. The first thrust aims to develop new marginal-based DPDS algorithms that improve upon the state-of-art in empirical evaluations. The tasks include: perform an in-depth study of the ?marginal-to-dataset? problem (how to synthesize a dataset when given a set of marginals); develop and evaluate new approaches for handling numerical attributes; and develop adaptive and automated techniques for selecting marginals so that dataset synthesized with them captures as much useful information from the input dataset as possible. The second thrust complements the empirical research in the first thrust, and aims to develop statistical theory for high dimensional marginal-based data synthesis algorithms, and also a general learning theory framework to evaluate the utility of synthetic data in downstream tasks. The two thrusts are highly complementary and support each other. The experimental study in Thrust 1 will provide insights and directions for theoretical studies in Thrust 2, which will help explain the experimental findings as well as guide additional experimental studies.

Project URL: http://www.cs.purdue.edu/homes/ninghui/projects/anonymization.html

## Privacy-Preserving Data Publishing

**PI: Ninghui Li**

**Current Students: Wahbeh Qadarji, Dong Su**

**Overview**

*Funded by National Science Foundation: TC: Small: Provably Private Microdata Publishing. 09/01/2011 - 08/31/2014.*

Data are a key resource in today's information age. The availability of data, however, often causes major privacy threats. Many data sharing scenarios require data to be anonymized for privacy protection. Most existing data anonymization techniques, however, satisfy only weak privacy notions that rely on particular assumptions about the adversaries, and provide inadequate protection. In recent years, the elegant notion of differential privacy has gradually been accepted as the privacy notion of choice for answering statistical queries. Most research on differential privacy, however, focuses on answering interactive queries, and there are several negative results on publishing microdata while satisfying differential privacy. Many data sharing scenarios, however, require sharing of microdata.

This project aims at bridging the gap between the elegant notion of differential privacy, and the practical difficulty of publishing microdata while preserving utility. Building on preliminary results by the PI, which have showed that random sampling plus ``safe" k-anonymization satisfies differential privacy, this project aims at advancing the state of the art of both the scientific understanding and the techniques for privacy-preserving microdata publishing. Research activities include developing (1) Practical anonymization methods that can be proven to satisfy differential privacy, while capable of handling high-dimensional data; (2) Relaxations of differential privacy that are more suitable for microdata publishing; (3) Privacy theory and techniques that are easily applied to a family of data sanitization algorithms called localized algorithms, enabling the usage of input perturbation techniques for provably private microdata publishing; (4) Privacy notions and techniques for publishing social network data and network trace data.

Project URL: http://www.cs.purdue.edu/homes/ninghui/projects/anonymization.html

## CRII: SaTC: Securing Smart Devices with AI-Powered mmWave Radar in New-Generation Wireless Networks

**PI: Tao Li**

**Overview**

New-generation networking systems, such as 5G/6G cellular and future Wi-Fi networks, are envisioned to connect billions of heterogeneous smart devices and enable high-speed and low-latency communications using millimeter wave (mmWave) technologies. The high frequency and short wavelength bring extraordinary sensing resolution which has enabled promising applications such as autonomous

driving, non-contact health monitoring, material detection, vibration sensing, etc. However, the research on security applications of mmWave sensing is still very much in its infancy. This project provides a comprehensive educational platform for curriculum development, graduate student training, undergraduate research engagement, and broadening the participation of women and underrepresented students in computing.

This project aims to develop, implement, and evaluate novel mmWave sensing techniques to secure smart devices such as smartphones, tablets, and laptops and the data therein in new-generation wireless networks. The project includes the following objectives:

- Developing mmLock, a user-leaving detection and identification system by analyzing the mmWave signal reflections from the user body to lock the lost/stolen device automatically and immediately against data theft when the user is unknowingly away from his/her device.
- Investigating mmUnlock, a liveness detection technique by checking consistency between the user's head movements extracted from the vision and mmWave data for more secure face authentication in smart devices.
- Implementing and evaluate the performance, security, and usability of above systems with commercial off-the-shelf (COTS) mmWave radars.

## Towards Machine-learnable Enhancing Framework for Local Differential Privacy

**PI: Wenhai Sun**

### Overview

The prevalence of data-centric applications demands collecting and analyzing the information that may contain sensitive data of users. While local differential privacy (LDP) can quantifiably control information leakage to protect user privacy, it is challenging to tackle a dynamic threat landscape where the attacker can manipulate the analytical results by exploiting the characteristics of LDP design. The project's novelties are to leverage machine intelligence along with other enabling techniques to understand the complex interplay between LDP privacy, security, and utility, and develop a machine-learnable LDP enhancing framework to address their tension. The success of the project will pave the way toward a future where ubiquitous machine intelligence can understand and attend to user privacy, security, and utility demands in various challenging scenarios. The project's broader significance and importance are broadening the participation of women and underrepresented students in STEM; encouraging interdisciplinary, cross-sector partnership, and technology transformation and incubation; and promoting diversity-and-equity-aware technology development.

This CAREER project aims to develop a novel AI-assisted privacy-enhancing framework that can better address the tension between LDP security, utility, and privacy. Generalizable theories and principles will be produced to guide the developed machine agent to sense the deployment environment and learn optimal responses to the observed adversarial actions, expected utility, and privacy goals. The agent will also be self-explainable for the actions it takes and can interact with both users and developers, thereby improving algorithmic transparency and accountability. In addition, human stakeholders will

get involved in the growing cycle of the system, allowing it to evolve over time. The project also aims to conduct research-informed educational activities that will strengthen cybersecurity-related education and mentoring programs in the home department while enhancing cybersecurity workforce training at the university level and beyond. Meanwhile, the project seeks to promote STEM careers by engaging in outreach activities with K-12 students with diverse backgrounds.

## In-Toto: Securing the Software Supply Chain

**PI: Santiago Torres-Arias**

**Overview**

The software development process is quite complex and involves a number of independent actors. Developers check source code into a version control system, the code is compiled into software at a build farm, and CI/CD systems run multiple tests to ensure the software's quality among a myriad of other operations. Finally, the software is packaged for distribution into a delivered product, to be consumed by end users. An attacker that is able to compromise any single step in the process can maliciously modify the software and harm any of the software's users.

To address these issues, we designed in-toto, a framework that cryptographically ensures the integrity of the software supply chain. in-toto grants the end user the ability to verify the software's supply chain from the project's inception to its deployment. We demonstrate in-toto's effectiveness on 30 software supply chain compromises that affected hundreds of million of users and showcase in-toto's usage over cloud-native, hybrid-cloud and cloud-agnostic applications. in-toto is integrated into products and open source projects that are used by millions of people daily.

**Representative Publications**

S. Torres-Arias, H. Afzali, T. K. Kuppusamy, R. Curtmola, andJ. Cappos, "In-toto: Providing farm-to-table guarantees for bits andbytes," inProc. 28th USENIX Sec., Santa Clara, CA, Aug. 2019

Project URL: https://in-toto.io

# SigStore: a Transparent Software Supply Chain Storage System

**PI: Santiago Torres-Arias**

**Overview**
Sigstore is a system to provide cross-ecosystem binary transparency, to register supply chain actors using federated identify management, and to allow software vendors to communicate software supply chain information between actors. Sigstore builds on existing transparent/auditable datastructures (e.g., transparency logs, transparency maps), as well as identification systems (e.g., OIDC), and supply chain metadata (e.g., in-toto, TUF, RPM signstures) to provide end-to-end verifiability to software consumers.

Project URL: https://sigstore.dev

# Towards Formal, Risk Aware Authorization

**PIs: David Yau, Adam Lee (University of Pittsburgh), Marianne Winslett (University of Illinois Urbana-Champaign)**

**Current Students: Naoki Tanaka, Heyu Xiong**

**Overview**
Traditional authorization decisions are black and white: a user either does or does not satisfy a given access policy. This rigidity is a handicap in our complex, dynamic, and unpredictable world. Proposals for risk-based access control address this problem by allocating principals a budget of risk tokens that can be used to buy access to sensitive resources. While this gives flexibility, pricing accesses to resources is non-trivial. Further, it is difficult to distinguish good and bad risk takers, and abandoning the formal proofs of authorization by traditional systems can lead to a lack of understanding of the actions taken in the system. The design of an adaptive access control system that is amenable to formal analysis thus remains an important open problem. To address this problem, we propose to develop a hybrid authorization approach that augments the strong formal guarantees of traditional attribute-based access control with more adaptive, risk-aware capabilities: risk-aware authorization (RAA).

To realize RAA, we willl design efficient, scalable methods to construct the best alternate proofs of authorization when a user cannot completely satisfy an ABAC access policy. We will automatically determine the minimum distance between these best alternate proofs and the traditional exact proofs for the desire access. We will use this minimum distance to determine the access price. We will use decision theory to justify the prices charged for risky accesses

and underpin the feedback loops in the system, which will use game theoretic methods to ensure that good risk takers are rewarded, bad risk takers are penalized, and everyone has an incentive to participate in the system. We will log the best alternate proofs to provide an audit trail justifying actions taken in the system. We will evaluate our results using case studies, formal proofs of correctness and optimality, and games that compare our automated strategies with those used by human players.

# A Novel Approach to Robust, Secured, and Cancellable Biometrics

**PIs: Xukai Zou, Yingzi Du, Scott Orr**

**Current Students:, Yan Sui, Tuo Lu, Zhi Zhou, Kai Yang**

## Overview

*Funding Source: Indiana University Center for Applied Cybersecurity Research*

Biometrics is to automatically identify or verify a person using physical, biological, and behavior characteristics, which include face, iris, fingerprints, hand geometry, voice, and etc. Compared to the traditional identification and verification methods (such as, some paper, plastic ID card, or password), biometrics is more convenient for users, reduces fraud, and can be more secure. Biometrics is becoming an important ally of security, intelligence, and law enforcement.

However, there are concerns about biometrics for daily life applications, such as security issues, privacy issues, standards, and etc. Among them, the biggest concern is the security of the biometric data. Unlike traditional identity methods, it is very hard, sometimes impossible, to re-issue a person's biometric data. If biometric data is obtained, for example compromised due to identity theft, the user will lose control over them forever and lose his/her identity.

Some researchers proposed to encrypt biometric data. They are using quite standard methods such as Advanced Encryption Standard (AES) and Public key cryptosystem RSA and cryptographic hash functions. The main issue related to them is key and key management, which has been studied independently from biometrics. As a result, there is a lack of research on the dependent relation between biometrics and encryption/integrity/key management and on comprehensive mechanisms involving authentication, encryption, data integrity, and key management.

Recently, some biometric researchers have proposed cancellable biometrics, which allows the system to re-issue the biometric for a user. The key idea of the cancellable biometrics is to distort the biometric image/signal/features before matching. The distortion parameters can be easily changed, which provides the cancelable nature of the scheme.

However, few if any have combined encryption and cancellable biometrics together to ensure the security of biometric data in storage, transmission, and identification. The simple and naïve approach is to put them together by designing a cancellable biometric method and applying encryption. This approach does not take consideration of the characteristics of biometrics and would not be applicable to real-life scenarios.

In this project, we propose a robust, secured, and cancellable biometrics method, which incorporates the encryption/key/key management into the cancellable biometric method design to provide the optimum solution. The PIs are experts in the field of biometrics, security, and network administration, which are essential for the success of this project.

# Evaluation of Clinical and Genomic Information Privacy Risks From Inference Attacks

**PIs: Xukai Zou, Jake Chen  Current Students: Huian Li**

## Overview

In this study, we will examine the quantitative relationships between clinical and genomic information disclosure and associated privacy risks due to inference attacks. For **inference attacks**, we refer to the inference of private personal identity and other personal information without the information owner's explicit consent or knowledge. In translational medical studies, identifiable personal information is usually anonymized and protected using a set of high-level guidelines. However, there is no explicit guarantee that such anonymization is performed to the best interests of research participants, especially with the increasing demand for open access of biobanks by researchers worldwide and, in some cases, patients themselves who are allowed to gain access to their own research results. Nor does there exist a method that can help researchers and biobank stakeholders gauge the risks for inference attacks, if the anonymized clinical database is compromised due to security leaks. Our specific aims are:

**Aim 1. Evaluate how clinical information, either disclosed through authorized or unauthorized access, may be used by inference attackers to reconstruct personal identity.**

- We will take actual metadata disclosed in actual cancer clinical studies in which the PI is currently involved in and perform simulated inference attacks to determine clinical information security (additional IRB approval will be sought).
- We will then broaden the research scope to include a survey of the literature-reported metadata collected in other clinical studies, in order to assess whether findings in our simulated attack have general applicability or not.

**Aim 2. Determine what sets of specific personal attributes and genomic variation loci may have a higher vulnerability for inference attacks if the security is compromised.**

- We will rank the set of common data attributes disclosed in clinical studies, based on their risk scores that we shall determine, based on our simulation results.
- We will also rank common genomic variation disclosed in Personal Genome Project (PGP) and the dbGAP database at NIH, to identify single nucleotide polymorphism (SNP) loci that are most discriminative of individuals.

The study has a **high potential impact** on future clinical and genomic data sharing/protection, which can be summarized as the following:

- The findings from this study can immediately benefit the data sharing design for hundreds of biobanking projects and clinical trial projects worldwide. These studies, which involve billions of dollars and millions of participants worldwide, will be able to re-examine possible information privacy vulnerability disclosed from this study and take actions to improve privacy protection.
- For biobanking or clinical study projects that require sharing of clinical data for research or individual use, the knowledge to be gained from this study (e.g., different privacy risk scores associated with each sets of clinical data or genotyping data) will help project investigators to make informed decisions that balance benefits and risks during information sharing.

# Secure Video Stream Framework for Dynamic and Anonymous Subscriber Groups

**PI: Xukai Zou**

**Current Students: Yan Sui, Kai Wang, Tuo Lu**

### Overview

*Funding Source: Cisco*

Secure video content distribution is a key aspect in the deployment of Telepresence Services and Video on Demand, two critical applications for the ecosystem targeted by Cisco products. Efficient mechanisms and systems need to be developed to guarantee confidentiality and controlled access to a broad range of broadcast video streams. At the same time, an effective framework for secure video content distribution should also guarantee subscribers' privileges to access video streams matching their respective subscription and on-demand requirements.

In this project, we will build, by employing an innovative approach called Access Control Polynomial (ACP), a Secure Video Stream Framework for dynamic and anonymous subscriber groups. The framework will effectively address the underlying challenges of secure video stream broadcasting and guaranteed access, anonymity, dynamicity, granularity, and scalability.

### Representative Publications

Y. Sui, F. Maino, Y. Guo, K. Wang, and X. Zou, An Efficient Time-bound Access Control Scheme for Dynamic Access Hierarchy, The Fifth International Conference on Mobile Ad-hoc and Sensor Networks (MSN'09) 14-16 December 2009, Wu Yi Mountain, China (Accepted).

K. Wang, X. Zou, and Y. Sui, A Multiple Secret Sharing Scheme based on Matrix Projection, COMPSAC'09, Seattle, WA, USA, July 20 –24, 2009, pp. 400-405.

# Trusted Medical Information System and Health Informatics

**PIs: Xukai Zou, Y. Dai (UTK)** **Current Students: Yan Sui, Kai Wang**

### Overview

*Funding Source: Department of Veterans Affairs*

In December of 2004 a US Marine is severely wounded during combat operations in Iraq. After receiving world class treatment at Bethesda Naval Hospital and the Indianapolis VA medical center, the patient is able to carry on a normal civilian life in Indianapolis. Several months later the veteran gets in an accident and is transported via medi-vac to a non-VA facility trauma center in Indianapolis for care. The provider looks up the patients data using the Indiana Health Information Exchange and the patient has a highly positive outcome. This outcome is only because critically important medical data was made available to the provider at the right time via a collaborative database between local hospitals. This scenario is only possible if VA hospitals can securely manage sharing of data between non VA health care facilities and themselves. The security schema the VA needs to meet this is a highly secure, manageable, portable, scalable, granular to the record & field level and most importantly cost effective security architecture.

It is with great enthusiasm we present the VISTALOCK security schema to the Department of Veterans Affairs. The scientists who have invented this technology are offering the Department of Veterans Affairs the opportunity to collaborate with them by implementing the already developed and proven technology across the VA Health Care domain. The VISTALOCK security architecture, using TEGO technology, is designed to be flexible and adaptable to support the security needs of VA and ALL of its national, regional and local affiliates.

VISTALOCK addresses four major security functions needed in collaborative data exchange and sharing, that is, Hierarchical Access Control (HAC), Secure Group Communication (SGC); Differential Access Control (DAC); Secure Dynamic Conferencing (SDC), enforces confidentiality, integrity, authentication, and fine tuned authorized access of patient records with granularity to the field and record level based on Cryptography and Key Management, and provides the capabilities of scalability, efficiency, dynamics, flexibility, and transparence.

The VISTALOCK security system is a bolt on security architecture that works in addition to the existing system(s) for which it protects, it will require no changes to the VISTA database repository and will act as a security gateway for all VISTA data traffic between the client and host. The VA will be able to apply best of breed technology to its security architecture, by providing modular and portable security services to the Vista/HealtheVET system. This enables the VA to continue full speed ahead with HealtheVET development as planned while still enabling secured collaborative data sharing capabilities to its architecture with external local health care facilities and practices.

### Representative Publications

X. Zou, Y.S. Dai, B. Doebbeling, and M. Qi, Dependability and Security in Medical Information System, the 12th International Conference on Human-Computer Interaction, 2007, Lecture Notes of Computer Science (LNCS), vol. 4553, pp.316-326.

Wang, Y. Sui, X. Zou, A. Durresi, and S. Fang, Pervasive and Trustworthy Healthcare. The First IEEE International Workshop on Bio Computing (BioCom'08), Okinawa, Japan, March 25 - 28, 2008, pp. 750–755.

# Autonomous Systems

## Adversarial Examples against Distributed Machine Learning Algorithms

**PIs: Saurabh Bagchi, David Inouye**

*Funding Source: U.S. Army Research Laboratory (ARL)*

### Overview

Adversarial examples (AEs) are images that can classifiers via introducing slight perturbations into original images. Recent work has shown that detecting AEs can be more effective than making the DNN robust against AEs. However, the state-of-the-art AE detection shows a high false positive rate, thereby rejecting a considerable fraction of normal images, and appears easy to bypass through reverse engineering attacks. To address this issue, we develop HAWKEYE, which is a separate classifier that analyzes the output layer of the DNN and detects AEs. Similar to prior work, HAWKEYE's AE detector utilizes a quantized version of an input image as a reference. However, instead of merely computing a simple statistic and then thresholding to detect AEs as in prior work, we train a separate simple classifier to distinguish the variation characteristics of the difference between the DNN output on an input image and the quantized reference image. By using a classifier, the detection rate is much higher, and thus we can cascade our AE detectors that are trained for different quantization step sizes to drastically reduce positive rate, while keeping the detection rate high.

Federated machine learning is a distributed learning approach that allows a global model to be trained across multiple decentralized client devices, e.g., IoT and edge devices. This approach offers privacy, security, and economic advantages to the participating clients by allowing for training using their local data, i.e., the model parameters are computed locally by the client devices and model parameter updates are shared with a central exchange server for iterative aggregation and consequent update of a global model. However, federated learning is subject to poisoning attacks abetted by the fact that no training examples are verified by a trustworthy authority. We present a typical federated learning scenario where the clients train their own local models on disjoint sets of data and periodically upload their local models to a parameter server for aggregation and download the aggregated (global) model. We show that it is possible for a malicious client to stealthily inflict an untargeted model poisoning attack, in contrast to a more traditional data poisoning attack, to disrupt training. Then, we present our defense technique, FLAIR,that can identify and remove the malicious clients to completely revive the training process and can be tuned to achieve a zero false positive detection rate over a wide range of usable hyperparameters.

# Addressing Safety and Security Challenges in ML-based AV Software Stack – Remote Operation Support and Balancing Trade-offs

**PIs: Yiheng Feng, Z. Morley Mao, University of Michigan**

*Funding Source: U.S. Department of Transportation*

### Overview

In this work, we propose a framework that enables safe and secure human remote operation when AV systems require support due to the inherent limitations of ML models used. We develop an approach that can effectively detect and mitigate potentially malicious remote human operators, and satisfy the real-time requirements of remote operation despite possibly variable network conditions impacting the communication channel between the AV system and the remote operator. Our solution will be demonstrated on the Mcity 2.0 testbed as a means to validate the proposed design in realistic settings.

Project URL: https://ccat.umtri.umich.edu/research/u-m/addressing-safety-and-security-challenges-in-ml-based-av-software-stack-remote-operation-support-and-balancing-trade-offs/

# Collaborative Research: CPS: Medium: Transforming Connected and Automated Transportation with Smart Networking, Cooperative Sensing, and Edge Computing

**PI: Yiheng Feng**

### Overview

Several recent technology trends are expected to transform ground transportation systems: much higherspeed wireless connectivity, improved on-vehicle and infrastructure sensing capabilities, and advances in machine learning. So far, most research and development efforts on connected and automated vehicles (CAVs) focused on individual technology. This leads to limited benefits, as communication, sensing, and computation functionalities are inherently tightly coupled. This project makes two foundational contributions in the ground transportation CPS domain. First, it advances the state-of-the-art of communication (focusing on the emerging 5G while embracing other wireless technologies), sensing (drastically scaling up cooperative raw sensor data sharing), and computation (applying distributed learning to a new domain) for CAV applications. Second, it proposes a run-time optimization framework that performs context-aware adaptation to meet application-specific requirements by jointly considering the networking, sensing, and computation dimensions. The project then applies the above innovations to support three CAV applications: enhancement of public service personnel's safety, alleviation of congestion at bottleneck areas, and protection of vulnerable road users (VRUs). In addition to simulationbased evaluations using CARMA, the team will prototype and evaluate the proposed work in Mcity, the world's first purpose-built proving ground for testing the performance of CAVs under controlled and realistic conditions.

Project URL: https://www.nsf.gov/awardsearch/showAward?AWD_ID=2038215&HistoricalAwards=false

# Object Recognition using Light Curve Inversion

**PI: Carolin Frueh, Vladimir Oliker, Emory University**

**Current Students: Alexander Burton, Liam Robinso**

### Overview

Light Curves are an easy and cheap way of collecting information on very distant objects, for which no resolved imaging can be acquired. For autonomous object recognition, light curve inversion is used to extract shape information from non-resolved imagery. The mathematical foundation is explored to find robust formulations for reliable shape inversion for realistic human-made objects, which include sharp edges and concavities.

### Representative Publications

S. Fan, C. Frueh, A Direct Light Curve Inversion Scheme and the Influence of Measurement Noise, Journal of Astronautical Science, submitted 2018

A. Buzzoni, D. Koschny, G. Drolshagen, E. Perozzi, O. Hainaut, S. Lemmens, G. Altavilla, I. Foppiani, J. Nomen, N. S?nchez-Ortiz, W. Marinello, G. Pizzetti, A. Soffiantini, S. Fan, C. Frueh, The observing campaign on the deep-space debris WT1190F as a test case for short-warning NEO impacts, Icarus, Vol. 304, pp. 4-8, doi:10.1016/j.icarus.2017.10.006, 2017

L. Robinson, C. Frueh, Light Curve Inversion for Reliable Shape-Reconstruction of Human-Made Space Objects, AAS/AIAA Astrodynamics Specialist Conference, Charlotte, North Carolina, August 7-11, 2022

# Observability for Autonomy and Sensor Network Design

**PI: Carolin Frueh**

**Current Students: Alex Friedman**

### Overview

A new formulation of the traditional concept of observability is used to allow setting up the optimal sensor network in order to retrieve the desired high-level information optimally. The connection to information measures is made. The formulation takes measurement noise, and sensor capabilities into account. The application is currently geared towards space objects but is not limited to that.

### Representative Publications

A. Friedman, C. Frueh, Determining Characteristics of Artificial Near-Earth Objects Using Observability Analysis, Acta Astronautica, Vol. 144, pp. 405-421, doi:10.1016/j.actaastro.2017.12.028, 2018

A. Friedman, C. Frueh, Observability Analysis Applied to Artificial Near-Earth Objects with Realistic Noise, Space Flight Mechanics Meeting, San Antonio, TX, Feb 2017

C. Frueh, Observability of Space Debris, AAS 15-576, Proc. AIAA/AAS Astrodynamic Specialist Conference, Vail, Colorado, August 2015

# Satellite Imaging using Compressed Sensing

**PI: Carolin Frueh**

**Overview**
Identification of satellites via imaging from non-resolved electro-optical measurements. This offers details about the satellite, which can otherwise not be determined.

**Representative Publications**
D. Kobayashi, C. Frueh, Reformulating Compressed Sensing to be used with Semi-Resolved Point Spread Function and Light Curves for Space Object Imaging: LEO. The Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference, September 27-30, 2022

D. Kobayashi, C. Frueh, Compressed Sensing for Satellite Characterization – Improved atmospheric modeling and sensing matrix realism, 8th European Conference on Space Debris, April 2021

# Hybrid UAM Model Checking

**PI: James Goppert**

**Current Students: Li-Yu Lin**

**Overview**
To analyze the complex concurrent interaction of aerial vehicles in the UAM system, we will employ model checking to verify logical properties in a finite state machine model. For example, if a no-fly zone is declared, all agents will leave the no-fly zone by a mandated deadline. Another example: Assuming random wind gusts of 20 knots, we will determine whether the planned UAM traffic can safely navigate without collision. We will verify the system using existing languages such as NuSMV50 and Spin51 by approximating the system as a Polyhedral-Invariant Hybrid Automaton (PIHA). Constructing a PIHA model allows verification of universal Computation Tree Logic (CTL) specifications. To construct a PIHA model, boundary certificates are not sufficient, since they do not compute a flow pipe. Flow pipes can be created through the creation of invariant sets around the reference trajectory and then sweeping the invariant set along it. Once flow pipes have been computed, the hybrid system can be converted to an approximate quotient transition system (AQTS) that can either over-approximate the reachable set of the system to prove safety or under-approximate it to prove reachability. The AQTS converts the hybrid model into a finite state machine that can be readily verified with existing model checkers.

Project URL: https://s2a2.ncat.edu

# Further Refinement and Integrated Platform for INDOT Traffic Management and Safety Toolset

**PI: Shu Hu** **Current Students: Mei Qiu**

*Funding Source: Joint Transportation Research Program (JTRP), Indiana Department of Transportation (INDOT)*

### Overview

This project addresses INDOT's need to use highway surveillance cameras to gather traffic information, such as flow rate, weaving data, and traffic anomaly detection. Specifically, TASI and INDOT will work on two aspects: (1) the expansion and refinement of the anomaly detection method being developed, (2) the development of a user-friendly system that integrates the software developed for lane-based flowrate detection, weaving analysis, and anomaly detection in the past and present, and other TASI and INDOT jointly developed traffic management tools in the near future.

The final system will be deployed to INDOT for daily operations.

### Representative Publications

Qiu, Mei and Lin, Wei and Chien, Stanley and Christopher, Lauren and Chen, Yaobin and Hu, Shu, Enhancing Vehicle Re-identification and Matching for Weaving Analysis, 2024 IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS) pp. 1-7

# Development of Verification and Validation System for Security Assessment of ECU Components

**PI: Younghyun Kim**

*Funding Source: Ministry of Trade, Industry and Energy, South Korea*

### Overview

This research aims to develop a standardized verification and validation (V&V) system for the automated verification of cybersecurity requirements in vehicle electronics components. The V&V system will autonomously execute verification procedures derived from the cybersecurity TARA analysis system, feed the results back into the TARA system, and be integrated into the security controller to monitor its internal state. Ultimately, this system aims to enable continuous security threat analysis and monitoring, ensuring robust and adaptive protection for vehicle electronics.

# Real Time CFD Mapping of Hazardous Airflow Around Bridge Infrastructure

**PI: John Mott** **Current Students: Kristoffer Borgen**

### Overview

Bridge inspections are an expensive and time-consuming process, varying significantly with the bridge's style, height, width, and length. Inspections create interruptions that interfere with bridge use, as the examination requires partial or total closure, causing traffic delays. Unmanned Aerial Systems (UAS) use has increased significantly over the past decade, including assistance and coordination during bridge inspections. However, the impact on the UAS from high winds and turbulent airflows induced by the bridge structure can decrease flight safety during inspections. Visualiza-tion of these hazards is difficult for UAS operators; therefore, a process to estimate the velocity and locations of these hazardous flows was created. The process begins by generating a simplified 3D model using the structural elements of a concrete and steel girder bridge based on the parameters and characteristics of the bridge. The model is then pro-cessed by a computational fluid dynamics (CFD) application that estimates the locations and velocities of the wind flows around the structure. Finally, the results are converted into a standard computer model file type that is either an augmented reality or computer application to display to assist the UAS operator.

Project URL: https://polytechnic.purdue.edu/advanced-aviation-analytics-institute-for-research

# AI-assisted Dynamic Adaptive Planning for Human-in-the-Loop Multi-Agent Systems

**PIs: Shaoshuai Mou, Dan Delaurentis, Xinhua Zhang, Joydeep Biswas**

**Current Students: Xuan Wang, Wanxin Jin, Paulo Heredia, Nick Schultz**

*Funding Source: Northrop Grumman*

### Overview

Topic 3 “Quantitative Dynamic Adaptive Planning” in Northrop Grumman's Research in Applications for Learning Machine Consortium (REALM). We aim to develop a AI-assisted Multi-Agent platform, which

- is able to provide distributed environment perception/situation awareness (on-board distributed fusion, AI-assisted object recognition)
- is able to perform real-time, dynamic and distributed control/management of assets, planning/ decision making/task assignment.
- is able to integrate human inputs in natural language/human gestures and provide feedback/ suggestions to human commander. (AI-assisted natural language processing, imitation learning, mixed-human-robot autonomy)
- is able to improve performance as time evolves (lifelong learning).

The platform will target at applications information gathering, search and rescue in response to nature disaster response.

The proposed research will be performed by a collaboration of four faculties from three different university with diverse research backgrounds. The team as a cohesive whole works together with Neta Ezer, Hasan A. Ghadialy and other Technical Leads from Northrop Grumman.

## Secure and Safe Assured Autonomy

**PIs: Shaoshuai Mou, Principal Investigator, Abdollah Homaifar, Co-Investigators, Daniel Delaurentis, Mark Costello, Ali Karimoddini, Inseok Hwang, Dengfeng Sun, Sam Coogan, Kyriakos G. Vamvoudakis, John C. Kelly, James Goppert, Yahya Kamalipour, Shaoshuai Mou, M. Nabil Mahmoud, Judy Hoffman, Ioannis A. Raptis, James Paduano, Bruce J Holmes, Damon Jenkins, Larry Datko**

**Overview**
Aviation's future will likely see the integration of a wide variety of Advanced Air Mobility (AAM) systems including Unmanned Aerial Systems (UAS) for cargo/delivery, personal air vehicles, and commercial Urban Air Mobility (UAM) vehicles. However, substantial challenges exist that could delay (and possibly prevent) these developments and thus research is needed in a variety of areas to leverage technologies in autonomy, Air Traffic Management (ATM), multi-redundant flight systems architectures, and advanced wireless connectivity like 5G to meet these challenges.
Project URL: https://uli.arc.nasa.gov/projects/10/

## Integrating Large-Scale Machine Learning and Edge Computing for Collaborative Autonomous Vehicles

**PIs: Dengfeng Sun, Joy Wang, Purdue ECE**

**Current Students: Bin Du**

**Overview**
The research objective of this project is to address the computational challenges in the innovative real-time and intelligent collaborative autonomous vehicles. A novel large-scale machine learning and edge computing framework is developed to integrate the emerging key computational techniques, including fast deep learning optimizations, asynchronous federated learning, cross domain deep learning model compression, hierarchical edge computing, and collaborative autonomous aerial and ground vehicles. Unlike most existing systems that perform big data analysis in central servers or clustering for offline learning, this project provides promising new directions to the real-time analysis of high-throughput sensor data by addressing the critical embedded device data analysis issues including efficiency, scalability, distributed computing, energy saving, and space reduction. The research project combines rigorous theoretical analysis and emerging application studies, and contributes to both academic research and potential commercialized products. Such unique capabilities enable new computational applications in a large number of research areas. It advances and thus extends the relationship between engineering innovation and computational analysis.

The goal of this ULI project is to develop new technologies and innovative operational concepts which will ensure safe, secure and robust integration of autonomous vehicles into Advanced Air Mobility-tailored transportation infrastructure. All this must be done while maintaining inter-operability with current civil air transportation systems and associated safety standards.

Project URL: https://www.nsf.gov/awardsearch/showAward?AWD_ID=1955890

## Secure and Safe Assured Autonomy

**PIs: Dengfeng Sun, Dan DeLaurentis, Inseok Hwang, Shaoshuai Mou, James Goppert**

**Current Students: Jiazhen Zhou**

### Overview

The team seeks to develop a novel integration of secure and safe autonomous systems used on unmanned Advanced Air Mobility (AAM) aircraft with the goal of advancing their technical readiness level and be ready for industry to consider using these technologies. The team intends to validate these systems with flight tests of multiple aircraft.

Project URL: https://web.ics.purdue.edu/~dsun/research.php

## Whitebox Testing, Debugging, and Repairing for Multi-module Autonomous Vehicles in Near-Collision Traffic Scenariosv

**PIs: Tianyi Zhang, Xiangyu Zhang**

**Current Students: Zhi Tu**

### Overview

This project seeks to develop principled algorithms and techniques for systematically testing, debugging, and repairing multi-module ADS to improve their safety and reliability. The core of our research is (1) a method for automated test-scenario construction that decouples high-level semantics and low-level details through a novel Domain Specific Language-based synthesis algorithm, (2) a search-based testing method that efficiently explores the enormous space of possible scenarios and identifies collision-inducing scenarios through a layered abstraction of multi-module autonomous systems and hierarchical optimization, and (3) a new adaptive debugging and repair technique that strategically diagnoses and fixes different kinds of safety bugs in different modules at different levels of granularity.

### Representative Publications

Deng, Yao, Xi Zheng, Mengshi Zhang, Guannan Lou, and Tianyi Zhang. "Scenario-based test reduction and prioritization for multi-module autonomous driving systems." In *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, pp. 82-93. 2022.

Lou, Guannan, Yao Deng, Xi Zheng, Mengshi Zhang, and Tianyi Zhang. "Testing of autonomous driving systems: where are we and where should we go?." In *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, pp. 31-43. 2022.

Deng, Y., Zheng, X., Zhang, T., Liu, H., Lou, G., Kim, M., & Chen, T. Y. (2022). A declarative metamorphic testing framework for autonomous driving. *IEEE Transactions on Software Engineering*, 49(4), 1964-1982.

Tu, Zhi, Liangkun Niu, Wei Fan, and Tianyi Zhang. "Multi-modal Traffic Scenario Generation for Autonomous Driving System Testing." P*roceedings of the ACM on Software Engineering* 2, no. FSE (2025): 1733-1756.

Y. Deng, Z. Tu, J. Yao, M. Zhang, T. Zhang and X. Zheng, "TARGET: Traffic Rule-Based Test Generation for Autonomous Driving via Validated LLM-Guided Knowledge Extraction," in *IEEE Transactions on Software Engineering*, vol. 51, no. 7, pp. 1950-1968, July 2025, doi: 10.1109/TSE.2025.3569086

# Cryptology and Rights Management

## Self-Healing Images

**PI: Hany Abdel-Khalik**

**Current Students: Arvind Sundaram, Dylan Adal**

**Overview**
The fragility of highly sensitive images has become a subject of recent concern to monitor and recover information in response to data corruption. The so-called self-healing image refers to the capability to restore lost information from a corrupted image, i.e., correcting pixelation or compression-induced imperfections. Furthermore, these methods serve as a self-authentication method in which the image can be confirmed as genuine by comparing top-level pixel information to embedded details; if there is a large discrepancy in the two, then it is clear that the images have been tampered with. The main purpose of this project is to construct a novel information processing algorithm that is applicable to multimodal data such as images and timeseries, wherein the true data is reinforced by self-healing capabilities, i.e., the ability to recover lost or damaged information.

The difficulty with typical methods is the ability to convey image clarity inversely decreases with the information embedded; one such method is adding information-carrying noise to the image, and many others utilize water-marking techniques. While the information potentially conveyed is very high, image clarity is quickly sacrificed in this procedure. The method proposed in this project allows for unrelated information-carrying components, e.g., timeseries-based information, another image's information, etc. to be embedded without compromising top-level behavior or visual acuity .

This method distinctly deviates from watermarking, which is overtly added to the data, and steganography, which can be discovered and reverse engineered, in that the proposed method is cryptographically secure by virtue of a randomized protocol that embeds information into the null space of a given image.

## The Garbled Computer: Towards Computing without Seeing

**PIs: Mike Atallah, Chris Clifton (Purdue), Qutaibah Malluhi and his colleagues from Qatar University**

**Overview**
This project aims at the development of a new secure computer model called the Garbled Computer (GC). An adversary observing the computations of a GC learns nothing about what it is doing, what data it is operating on (whether inputs or intermediate values), and the outputs it is producing. The GC achieves, using a single general approach, the multiple goals of software obfuscation, tamper-proofing, data confidentiality and data integrity. The GC enables execution, on untrusted platforms, of trusted and confidential code whose inputs and outputs are sensitive. For example, it can enable the utilization of Amazon cloud services without revealing to Amazon the nature of the computation or the data, and without requiring Amazon to change the operation of its cloud services (i.e., use standard off-the-shelf services).

# Password Hashing Algorithms

**PI: Jeremiah Blocki**

**Current Students: Ben Harsha, Samson Zhou, Seunghoon Lee**

## Overview

In the last few years breaches at organizations like Yahoo!, Dropbox, Lastpass, AshleyMadison and Adult FriendFinder have exposed over a billion user passwords to offline attacks. Password hashing algorithms are a critical last line of defense against an offline attacker who has stolen password hash values from an authentication server. A attacker who has stolen a user's password hash value can attempt to crack each user's password offline by comparing the hashes of likely password guesses with the stolen hash value. Because the attacker can check each guess offline it is no longer possible to lockout the adversary after several incorrect guesses. The attacker is limited only by the cost of computing the hash function. Offline attacks are increasingly commonplace and dangerous due to weak password selection and improved cracking hardware e.g., the Antminer S9, currently available on Amazon.com for around $3,000 (USD), is capable of computing 14 trillion SHA256 hashes/second. When LastPass was breached they were using PBKDF2, a slow password hashing algorithm which iteratively computes SHA256 100,000 times. Thus, a LastPass attacker could potentially check 140 million password guesses per second on the Antminer S9. By comparison, 70 million guesses suffice to crack most user passwords (e.g., see <a href="”https://eprint.iacr.org/2016/153”>empirical frequency data for Yahoo! passwords</a>). There is a clear need to develop secure (moderately expensive) password hashing algorithms so that it is economically infeasible for an offline adversary to check millions of password guesses.

Recognizing this clear need researchers recently organized the Password Hashing Competition (PHC) to encourage the development of better password hashing algorithms. A secure password hashing algorithm should be: 1) quickly computable (e.g., $< 1$ second) on a PC, 2) prohibitively expensive for an adversary to evaluate millions of times even using customized hardware e.g., an Application Specific Integrated Circuit (ASIC) like the Antminer S9. Memory hard functions (MHFs), functions whose computation require a large amount of memory, are a promising primitive to achieving both goals. Memory is expensive even on an ASIC, and the Area x Time (AT) complexity of computing an ideal MHF scales with $n^2$, where $n$ is the running time on a PC. Thus, an MHF allows us to rapidly increase costs. By contrast, password hash functions like PBKDF2 or BCRYPT require minimal memory to compute and the AT complexity only scales with $n$. Consequently, PBKDF2 and BCRYPT can both be evaluated on an ASIC at minimal cost. Data-Independent Memory Hard Functions (iMHFs) are an important variant of MHFs in the context of password hashing due to their resistance to side-channel attacks.

## Representative Publications

Depth-Robust Graphs and Their Cumulative Memory Complexity. with Joel Alwen and Krzysztof Pietrzak. EUROCRYPT 2017)..

Towards Practical Attacks on Argon2i and Balloon Hashing. with Joel Alwen. EuroS&P 2017

Efficiently Computing Data Independent Memory Hard Functions. with Joel Alwen. CRYPTO 2016.

On the Computational Complexity of Minimal Cumulative Cost Graph Pebbling. with Samson Zhou. Financial Crypto 2018.

Sustained Space Complexity. with with Joel Alwen and Krzysztof Pietrzak. EUROCRYPT 2018.

Bandwidth-Hard Functions: Reductions and Lower Bounds. with Ling Ren and Samson Zhou. 25th ACM Conference on Computer and Communications Security.

Practical Graphs for Optimal Side-Channel Resistant Memory-Hard Functions. with Joel Alwen and Ben Harsha 24th ACM Conference on Computer and Communications Security

## Decentralized Anonymous Credentials from Blockchains

**PI: Christina Garman**

**Overview**
Anonymous credentials provide a powerful tool for making assertions about identity while maintaining privacy. However, a limitation of today's anonymous credential systems is the need for a trusted credential issuer — which is both a single point of failure and a target for compromise. Furthermore, the need for such a trusted issuer can make it challenging to deploy credential systems in practice, particularly in the ad hoc network setting (e.g., anonymous peer-to-peer networks) where no single party can be trusted with this responsibility.

This work will explore expanding on prior work in decentralized anonymous credentials, where, using a decentralized ledger (like a blockchain) and standard cryptographic primitives, we built an anonymous credentials system that allows users to make flexible identity assertions with strong privacy guarantees. We will also explore various applications of these credentials, including Tor.

## Flexible Anonymous Credentials from zk-SNARKs and Existing Identity Infrastructure

PI: Christina Garman

**Current Students: Jacob White**

**Overview**
Anonymous credentials provide a powerful tool for making assertions about identity while maintaining privacy. While the subject of extensive academic work, anonymous credentials have thus far seen little deployment. In large part, this is because most existing systems are designed with a number of assumptions about identity that, in fact, are not actually true in practice, rendering them difficult to

actually deploy. For example, existing systems assume that: there is a single issuer for a given identity property; there exist reputable authorities that are able (and more importantly willing) to be responsible for holding keys, verifying identity properties, and issuing credentials; all use cases and attributes needed for a credential are known in advance and are only simple statements ("my age is"); and the set of authorities for a given identity attribute or credential can be enumerated at the time we instantiate the system.

This work will explore expanding on prior work in decentralized anonymous credentials, where, using a decentralized ledger (like a blockchain) and standard cryptographic primitives, we built an anonymous credentials system that allows users to make flexible identity assertions with strong privacy guarantees. Our aim is to now build a flexible, issuer-agnostic anonymous credential toolkit which allows for the representation of complex identity statements while still supporting the traditional properties afforded by previous anonymous credential systems. Additionally, the system will allow for widespread adoption by supporting the use of existing identity infrastructure (such as passports or drivers licenses) and enabling organizational agility, as both issuers and those accepting credentials do not need to work together to support new use cases.  We will also explore various real world applications of these credentials.

## Privacy Preserving Software Bill of Materials

**PI: Christina Garman**

**Current Students: Arushi Arora**

### Overview

**Funding Source: Idaho National Laboratory**

Modern software is generally composed of a number of different libraries and subcomponents. Recent severe vulnerabilities, such as Log4j, have highlighted the necessity of understanding and cataloging all of the components in a system or piece of software. A Software Bill of Materials (or SBoM) is designed to do just this, providing a formal record that enumerates all the components of a given piece of software, along with their respective relationships. While there have been a number of initial recommendations for and deployments of SBoMs, we observe that there is significant potential for privacy concerns in such deployments. SBoMs may relate to proprietary software or subcomponents, and a publicly-accessible SBoM could potentially aid an attacker in discovering vulnerable products in the case of a library exploit. As such, we seek to investigate if a software owner can conceal the contents of an SBoM from a prospective user, while still allowing the user to benefit from the full SBoM feature set. To do this, we introduce the notion of a privacy-preserving SBoM system, and define its desired security properties and necessary components. We then provide a concrete instantiation of such a scheme, building off of ideas from the Private Set Intersection space. To demonstrate our scheme's practicality, we provide a full end-to-end implementation, including an integration with a real world SBoM system, as well as a series of benchmarks and discussion of real word deployment concerns.

# Removing the Human Element: Securing Deployed Cryptographic Systems through the use of Cryptographic Automation

**PI: Christina Garman**

## Overview

*Funding Source: National Science Foundation (NSF)*

Cryptography has shown itself to be invaluable in everyday life, especially as more and more devices and interactions are moving to the online world. Whether it is browsing the web, making a purchase, or sending a message to a friend, cryptography is everywhere. Despite the fact that users (often unknowingly) rely on the security of systems that use cryptography, recent years have seen a number of serious vulnerabilities in the cryptographic pieces of systems, some with large consequences. These have been caused by various problems, including poor designs, difficulty of implementation, and use (or misuse) of (in)secure primitives. There is a common denominator in all of these problems: the human element. Many of the errors that are found when analyzing these insecure systems could have been prevented if both designers and software engineers had better tools to help them navigate the complex cryptographic space. Cryptographic automation is a relatively new and promising area that is designed to help solve many of these issues and make developing secure systems far easier and less error-prone, even for a non-expert. This project focuses on removing the human element from the deployment and analysis of cryptographic systems. Through the use of cryptographic automation and the development of tools, the project's aim is to make it easier to design and securely deploy new and complex cryptographic systems while preventing insecurities from occurring in such systems. Additionally, the project contains an education plan designed to help make cryptography more accessible to a broader audience. The creation of the Midwest Women in Computer Security Workshop, as well as the project's goal to not just develop but also disseminate tools, will allow more students of all ages, and more software engineers, to explore cryptography and computer security, instead of being intimidated or afraid of it.

The project has three main thrusts. The core of the project centers around the first thrust of building tools to aid in the deployment of complex cryptography. This will principally focus on automating the end-to-end development of zero-knowledge proof code, from expressing the proof statement to realizing the implementation, with additional applications to anonymous credentials. The second thrust focuses on automating the discovery of cryptographic vulnerabilities in applications that use zkSNARKs, a popular zero-knowledge proof instantiation. This thrust will leverage fuzzing to help both programmers and end users detect inconsistencies and errors in existing, already deployed zkSNARK circuits and applications. The third thrust works to automate the discovery and identification of modern cryptographic algorithms and techniques in both traditional as well as heavily obfuscated binaries, through a novel combination of various dynamic analysis and machine-learning based approaches. If successful, the combination of these three thrusts will, for expert and non-expert developers alike, make it both easier to discover the use of cryptography and potentially vulnerable algorithms in existing systems as well as design and securely deploy new and complex cryptographic systems while preventing these insecurities from happening.

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

**PI: Christina Garman**

**Current Students: Jacob White**

## Overview

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 these concerns. However, existing schemes do not handle the realities of deployment or the complexities of real-world identity. Instead, they implicitly make assumptions such as there being an issuing authority for anonymous credentials that, for real applications, requires the local department of motor vehicles to issue sophisticated cryptographic tokens to show users are over 18. In reality, there are 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.

## Representative Publications

M. Rosenberg, J. White, C. Garman and I. Miers, "zk-creds: Flexible Anonymous Credentials from zkSNARKs and Existing Identity Infrastructure," in 2023 2023 IEEE Symposium on Security and Privacy (SP) (SP), San Francisco, CA, US, 2023 pp. 1882-1900.
doi: 10.1109/SP46215.2023.00108

Project URL: https://github.com/rozbb/zkcreds-rs

# SaTC: Frontiers: Center for Distributed Confidential Computing (CDCC)

**PI: Ninghui Li**

**Current Students: Weicheng Wang**

*Funding Source: National Science Foundation (NSF)*

## Overview

Advances in AI and big data analytics rely on data sharing, which can be impeded by privacy concerns. Most challenging in privacy protection is protection of data-in-use, since even encrypted data needs to be decrypted before it can be utilized, thereby exposing data content to unauthorized parties. A practical and scalable solution to the challenge will transform computing, enabling unprecedented capabilities such as confidential outsourcing, trusted computing services, and confidential or privacy-preserving collaboration. In quest of such a holy grail of data protection, this frontier project establishes multi-institution and multi-disciplinary Center for Distributed Confidential Computing (CDCC) to create a research, education, knowledge transfer and workforce development environment that enables scalable, practical, verifiable and usable data-in-use protection based upon Trusted Execution Environments (TEE) on cloud and edge systems.

CDCC focuses on four building block thrusts fundamental to distributed confidential computing (DCC), regardless of specific TEE hardware, including assurance of TEE code, assurance of TEE nodes, assurance of a TEE workflow and assurance for the stakeholder. The first thrust leads to an open ecosystem for TEE code certification, not relying on any trusted party but on a trustworthy application store whose certification operations are public, accountable and verifiable. The second thrust aims to develop novel dynamic data-use policy models and enforcement mechanisms for scalable trust management and data control on the TEE nodes running certified code. The third thrust focuses on ensuring protection of the computational workflow built on TEE nodes and the last thrust studies the stakeholder's preference and expectations to guide the design of DCC technologies and ensure their usability. On top of these building blocks, the center explores various transformative applications (e.g., confidential distributed AI supports for healthcare) to be enabled. CDCC also has a number of efforts for outreach (development of a massive open online course, industry collaboration, etc.) and for broadening participation (security and privacy lab for attracting minority students, joint summer schools and others).

# Cyber-Physical Systems

## Covert Cognizance

**PI: Hany Abdel-Khalik** **Current Students: Arvind Sundaram**

**Current Students: Arvind Sundaram**

*Funding Source: Department of Energy (DOE), U.S. Army Research Laboratory (ARL)*

### Overview

The Covert Cognizance (C2) paradigm has been developed in response to the growing frequency of cyberattacks, specifically insider threats and state-sponsored advanced persistent campaigns. In such scenarios, adversaries possess the technical know-how and financial resources to bypass IT-based security defenses such as firewalls, passwords, biometrics etc. as well as model-based defenses that rely on a physical model or a digital twin of the cyberphysical system (CPS). As such, there is a need for a human-free and deterministic solution to cybersecurity, a last line of defense, when all IT and OT defenses have been compromised.

C2 seeks to embed self-awareness in CPS at level 0 of the traditional Purdue model, i.e., at the level of the sensors and actuators which are the most critical and sensitive components in a CPS. Departing from protocols that shield the data flowing through these components by encasing it in a (penetrable) shell, C2 directly fingerprints the data at its source (level 0) by embedding information, called C2 parameters, about the system's operational history in the data. This effectively makes the components cognizant or "aware" of each other. In traditional models, these components are loosely coupled with each other through the physics of the CPS which are well-understood and may be learned and exploited using AI/ML. In C2-enabled CPS, these systems are tightly coupled through imperceptible perturbations that carry the C2 parameters. Any falsification of the data will necessarily destroy the presence of the C2 parameters, thus serving as a deterministic tool to intrusion detection.

Another key value of C2 is that the C2 parameters can also carry recovery information to prevent downtime of systems during cyberattacks. In essence, C2-enabled CPS are a system of systems that can detect cyberattacks deterministically and instantly self-heal to nullify their effect without needing human intervention. All these capabilities are endowed in a manner that cannot be reverse-engineered with security guarantees based on the Vernam-cipher/one-time-pad gold standard along with operational guarantees based on the criterion of zero-impact on system optimality.

### Representative Publications

1 Arvind Sundaram, and Hany S. Abdel-Khalik, "Validation of Covert Cognizance Active Defenses", Nuclear Science and Engineering, April 2021.

2 Arvind Sundaram, and Hany S. Abdel-Khalik, "Covert Cognizance: A Novel Predictive Modeling Paradigm", Nuclear Technology, February 2021.

3 Arvind Sundaram, Hany S. Abdel-Khalik, and Oussama Ashy, "A data analytical approach for assessing the efficacy of Operational Technology active defenses against insider threats", Progress in Nuclear Energy, June 2020.

4 Arvind Sundaram, and Hany S. Abdel-Khalik, "Developing Covert Cognizance for Industrial Control Systems", International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering, M&C 2021, October 2021.

5 Arvind Sundara, and Hany S. Abdel-Khalik, "Exploratory Study into the Effectiveness of Active Monitoring Techniqus", Transactions of the American Nuclear Society 2019 Winter Meeting, December 2019.

## Securing Embedded Devices by Enforcing Lowest Privilege Execution

**PIs: Saurabh Bagchi, Mathias Payer, Abraham Clements**

*Funding Source: Sandia National Laboratories*

### Overview

With more than 9 billion embedded processors in use today, the number of embedded devices has surpassed the number of humans. With the rise of the "Internet of Things" (IoT), the number of embedded devices and their connectivity are exploding. These "things" include fitness trackers, smart light bulbs, smart thermostats, home assistants, utility smart meters, and smart locks. The increasing network connectivity coupled with the ubiquity of these devices makes securing IoT systems a critical task. Evidence of the dangers of insecure IoT systems abounds. For example, in 2016, hijacked smart devices like CCTV cameras and digital video recorders launched the largest distributed denial of service attack to date.

Many of these devices are low cost with software running directly on the hardware, known as "bare-metal systems." The application runs as privileged, low-level software with direct access to the resources of the microcontroller (μC) and its peripherals. Unlike desktop systems, there are no intervening operating system software layers to control access to the resources in a secure manner. Making matters worse, embedded systems largely lack protection against code injection, control-flow hijacking, and data corruption attacks.

We are improving the security for bare-metal embedded and IoT systems through the design of a privilege overlay that restricts the privileges and capabilities of different regions of the application to the lowest necessary to perform intended operations. Our innovations instantiated in a system called ACES can be applied without needing application modification and with limited user annotations, to indicate what denotes security-critical operations, thus easing the application of ACES to legacy embedded applications. We have achieved this through three interlocking tasks:

- New static and dynamic analyses to identify security and functionality characteristics of each part of the application;
- New runtime techniques that enforce the desired security properties while minimizing the performance impact; and
- New security metrics and benchmarks that accurately measure the security and performance impacts of defense mechanisms for embedded systems.

We demonstrate the benefits of ACES through realistic applications developed in five domains on actual

hardware—smart homes, wearables, smart cities, transportation, and industrial control systems.

### Representative Publications

Abraham A. Clements (Purdue & Sandia), Eric Gustafson (UCSB), Tobias Scharnowski (Ruhr University Bochum), Paul Grosen (UCSB), David Fritz (Sandia), Christopher Kruegel (UCSB), Giovanni Vigna (UCSB), Saurabh Bagchi, and Mathias Payer (EPFL), "HALucinator: Firmware Re-hosting Through Abstraction Layer Emulation," Accepted to appear at the 29th USENIX Security Symposium (Usenix Sec), pp. 1-18, Aug 12-14, 2020, Boston, MA.

Naif Saleh Almakhdhub (Purdue and King Saud University), Abraham A Clements (Sandia National Labs), Saurabh Bagchi, and Mathias Payer (EPFL), "μRAI: Return Address Integrity for Embedded Systems," At the Network and Distributed System Security Symposium (NDSS), pp. 1–18, February 23-26, 2020, San Diego, CA.

Naif Almakhdhub, Abraham Clements, Mathias Payer and Saurabh Bagchi, "BenchIoT: A benchmark for the things in the Internet of Things," At the 49th IEEE/IFIP International Symposium on Dependable Systems and Networks (DSN), pp. 234-246, June 24-27, 2019, Portland, OR.

Abraham A. Clements, Naif S. Almakhdhub, Saurabh Bagchi, and Mathias Payer, "ACES: Automatic Compartments for Embedded Systems," In Proceedings of the 27th USENIX Security Symposium (USENIX Sec '18), pp. 65–82, Aug 15–17, 2018, Baltimore, MD.

Abraham A Clements, Naif Saleh Almakhdhub, Khaled Saab, Prashast Srivastava, Jinkyu Koo, Saurabh Bagchi, and Mathias Payer, "Protecting Bare-metal Embedded Systems with Privilege Overlays," In Proceedings of the IEEE International Symposium on Security and Privacy (Oakland), pp. 289-303, May 22-26, 2017, San Jose, California.

Project URL: https://engineering.purdue.edu/dcsl/project-new/project-1/

## Characterization of soft dielectric breakdown in GaN MOSHEMTs

**PIs: Peter Bermel, Peide Ye, Ali Shakouri, Justin Weibel**

**Current Students: David Kortge, Jinhyun Noh, Xiao Lyu**

### Overview

Power electronics convert alternating to direct current, and play a critical role in cyberphysical system infrastructure, from personal transportation to grid-level battery storage. Current state-of-the-art power electronics are gallium nitride (GaN) high electron mobility transistors (HEMTs). GaN HEMTs support higher operating voltages, allowing for highly efficient operations, but are subject to high power drains in the off state through gate current leakage. GaN MOSHEMTs are an attractive successor to GaN HEMTs since the addition of a gate oxide significantly reduces the gate leakage. However, before MOSHEMTs can be widely adopted, their reliability and corresponding modes of degradation must be well understood and predictable. In this work, soft breakdown of the gate dielectric in GaN MOSHEMTs is characterized and

modeled using a constant voltage source for degradation. This will enable work at higher levels to design devices and systems with well-characterized failure rates for secure cyberphysical system applications.
Representative Publications

Noh, Jinhyun, Sami Alajlouni, Marko J. Tadjer, James C. Culbertson, Hagyoul Bae, Mengwei Si, Hong Zhou, Peter A. Bermel, Ali Shakouri, and D. Ye Peide. "High Performance Beta-Gallium Oxide Nano-Membrane Field Effect Transistors on a High Thermal Conductivity Diamond Substrate." *IEEE Journal of the Electron Devices Society* 7 (2019): 914-918.

## Compositional IoT Safety and Security in Physical Spaces

**PI: Berkay Celik**

**Overview**
An IoT system is composed of multiple individual components, with each component containing a set of sensors and actuators governed by a control program. The inevitable integration of many individual system components programmed independently into IoT systems has brought new challenges that require urgent attention. The main challenge is to produce proofs of correctness that ensure the composite behavior of IoT devices in physical spaces—the environment in which sensors/actuators operate—adheres to desired safety and security policies.

Celik's project integrates research activities aimed at designing and developing algorithms and tools that formally produce the composite behavior of an IoT system and a rigorous foundation for reasoning about an IoT environment's global safety and security.

The specific goals of the project are divided into three research thrusts. The first thrust focuses on constructing a novel composite model by unifying the behavior of individual system components through a combination of static analysis and system identification techniques to represent an IoT system's global behavior. The second thrust aims to establish a rigorous foundation for identifying physical behavior-based policies and developing formal analysis techniques that ensure an IoT system adheres to safety and security policies. The last thrust seeks to establish a series of techniques to make model construction and policy validation scalable and exhaustive in diverse IoT systems.

## Bringing Fuzzing to the Cyber-Physical World

**PIs: Berkay Celik, Dave Tian, Dongyan Xu, Antonio Bianchi**

**Overview**
This grant, named Bringing Fuzzing to the Cyber Physical World, will investigate how to use fuzzing techniques to discover bugs and vulnerabilities in Cyber-Physical Systems.

Project URL: https://pursec.cs.purdue.edu/

# Security of Autonomous Vehicles

**PIs: Berkay Celik, Antonio Bianchi**

*Funding Source: National Science Foundation (NSF), Office of Naval Research (ONR)*

**Overview**
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.

**Representative Publications**
Discovering Adversarial Driving Maneuvers against Autonomous Vehicles
https://www.usenix.org/conference/usenixsecurity23/presentation/song

Project URL: https://beerkay.github.io/

# Formal methods and Fuzzing for Security in Internet of Things, Embedded Systems,Real-time Operating Systems, and General Software

**PIs: Jamie Davis, Aravind Machiry (ECE)**

**Current Students: Paschal Amusuo (PhD student @ECE), Sid Muralee (PhD student @ECE), Ritvik Tanksalkar (PhD student @ECE)**

*Funding Source: Funding being sought, Rolls-Royce*

**Overview**
Embedded software makes the world go 'round. Static and dynamic analysis of embedded software is a necessary capability for high-assurance software engineering (e.g. IEC 61508, ISO 26262). These capabilities facilitate many security tasks, e.g., vulnerability detection and repair, and reverse engineering.

This project evaluates the security of embedded software applications as well as the infrastructure on which they depend, such as embedded network stacks (lwIP etc.) and real-time operating systems (RTOSes like FreeRTOS). We are evaluating static analysis options and identifying shortcomings. We are performing vulnerability analysis to identify common weaknesses across vendors. We are working on automated rehosting to apply state-of-the-art dynamic analysis techniques (e.g. fuzzing) in a UNIX environment.

### Representative Publications

*Usage and Effectiveness of Static Analysis in Open-Source Embedded Software: CodeQL Finds Hundreds of Defects. Shen, Yuan, Pillai, Zhang, Davis, and Machiry. arXiv 2024.*

*A Unified Taxonomy and Evaluation of IoT Security Guidelines*. Chen, Anandayuvaraj, **Davis**, and Rahaman. arXiv 2023.

*Towards Rehosting Embedded Applications as Linux Applications*. Srinivasan, Tanksalkar, Amusuo, **Davis**, and Machiry. Proceedings of the 53rd Annual IEEE/IFIP International Conference on Dependable Systems and Networks — Disrupt track (DSN-Disrupt) 2023.

*Systematically Detecting Packet Validation Vulnerabilities in Embedded Network Stacks*. Amusuo, Méndez, Xu, Machiry, and **Davis**. Proceedings of the 38th IEEE/ACM International Conference on Automated Software Engineering (ASE'23) 2023.

*Towards Automated Identification of Layering Violations in Embedded Applications (WIP).* Shen, **Davis**, and Machiry. Proceedings of the 24th ACM SIGPLAN/SIGBED International Conference on Languages, Compilers, and Tools for Embedded Systems (LCTES) 2023.

*"If security is required": Engineering and Security Practices for Machine Learning-based IoT Devices*. Gopalakrishna, Anandayuvaraj, Detti, Bland, Rahaman, and **Davis**. Proceedings of the 4th International Workshop on Software Engineering Research & Practices for the Internet of Things (SERP4IoT) 2022.

*LEMIX: Enabling Testing of Embedded Applications as Linux Applications*. Tanksalkar, Muralee, Danduri, Amusuo, Bianchi, **Davis**, and Machiry.
Proceedings of the 34th USENIX Security Symposium (SECURITY) 2025.

*Reactive Bottom-Up Testing*. Muralee, Cherupattamoolayil, **Davis**, Bianchi, and Machiry. arXiv 2025

*Do Unit Proofs Work? An Empirical Study of Compositional Bounded Model Checking for Memory Safety Verification*. Amusuo, Cochell, Le Lievre, Patil, Machiry, and **Davis**.
Proceedings of the 48th IEEE/ACM International Conference on Software Engineering (ICSE) 2026.

*ZTD-JAVA: Mitigating Software Supply Chain Vulnerabilities via Zero-Trust Dependencies*. Amusuo, Robinson, Singla, Peng, Machiry, Torres-Arias, Simon, and **Davis**.
Proceedings of the ACM/IEEE 47th International Conference on Software Engineering (ICSE) 2025.

*A Unit Proofing Framework for Code-level Verification: A Research Agenda*. Amusuo, Patil, Cochell, Le

Lievre, and **Davis**. Proceedings of the 47th IEEE/ACM International Conference on Software Engineering - New Ideas and Emerging Results Track (ICSE-NIER) 2025.

*FalseCrashReducer: Mitigating False Positive Crashes in OSS-Fuzz-Gen Using Agentic AI*. Amusuo, Liu, Calvo, Metzman, Chang, and **Davis**. arXiv 2025.

## Promoting Inter- and Intra-Organizational Learning from Software Failures: Towards a Failure-Aware Software Development Lifecycle

**PI: Jamie Davis** **Current Students: Dharun Anandayuvaraj, PhD student**

*Funding Source: Funding being sought*

### Overview

The goal of this project is to help software engineers incorporate the lessons learned from prior failures throughout the software development lifecycle. Since all engineered systems fail, one fundamental theorem of software engineering is to learn from failures to mitigate their recurrence. Although this theorem has been recommended for software engineers by standards bodies (e.g., the ISO) and organizations such as Google (e.g., the SRE book), it has received little critical examination. We are conducting foundational empirical research to understand current and best practices related to software failure feedback. We will use this knowledge to develop and evaluate innovations in failure-aware software development processes.

Some software engineering practices, such as postmortems and retrospectives, render failure knowledge into lessons learned. These lessons may be incorporated into other engineering processes and artifacts, e.g., design reviews and style guides. The problem is that we lack basic empirical data to inform the application of this feedback mechanism. For example, we do not know: (1) What are best practices in failure knowledge collection and sharing? ; nor (2) How well is failure knowledge leveraged in the software development lifecycle? ; nor (3) How and to what extent do engineers study failures in other organizations' products to inform their own work? ; nor (4) What is the cost of failure analysis vs. the benefit in future failure elimination? Answering such questions could be transformative.

Our research goal is to collect these basic empirical data. To do so, we will develop tooling and conduct human-subjects work to facilitate inter- and intra-organizational learning. The expected outcomes are knowledge about current and best practices, and preliminary evaluations of innovations in failure-aware development practices. If successful, our results will enable engineers to analyze failures and apply lessons throughout the software development lifecycle; and enable engineering decision-makers to determine which failures to focus on and how to assess the cost/benefit trade-offs in their context.

### Representative Publications

*An Empirical Study on Using Large Language Models to Analyze Software Supply Chain Security Failures*. Singla, Anandayuvaraj, Kalu, Schorlemmer, and **Davis**. Proceedings of the 2nd ACM Workshop on Software Supply Chain Offensive Research and Ecosystem Defenses (SCORED) 2023.

*Incorporating Failure Knowledge into Design Decisions for IoT Systems: A Controlled Experiment on Novices*. Anandayuvaraj, Thulluri, Figueroa, Shandilya, and **Davis**. 5th International Workshop on Software Engineering Research & Practices for the Internet of Things (SERP4IoT 2023) 2023.

*Reflecting on Recurring Failures in IoT Development*. Anandayuvaraj and **Davis**. Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering: New Ideas and Emerging Results track (ASE-NIER) 2022.

*A Unified Taxonomy and Evaluation of IoT Security Guidelines*. Chen, Anandayuvaraj, **Davis**, and Rahaman. arXiv 2023.

*FAIL: Analyzing Software Failures from the News Using LLMs*. Anandayuvaraj, Campbell, Tewari, and **Davis**. Proceedings of the 39th IEEE/ACM International Conference on Automated Software Engineering (ASE) 2024.

*Learning From Software Failures: A Case Study at a National Space Research Center*. Anandayuvaraj, Hammadeh, Lund, Holloway, and **Davis**. arXiv 2025.

## CAREER: Securing Next-Generation Transportation Infrastructure: A Traffic Engineering Perspective

**PI: Yiheng Feng**

### Overview

This project establishes an integrated research and education plan to investigate the cybersecurity risk in next-generation transportation infrastructure, which is envisioned to be equipped with advanced sensors, communication systems, and edge computing capabilities to provide safer and more efficient traffic management strategies. Despite the benefits, these components are increasingly exposed to cyberattacks. To secure transportation infrastructure in cyberspace, this project conducts security analysis, recommends defense solutions, and evaluates system performance under various scenarios. The results inform public agencies for developing standards and policies, and prioritizing deployment of limited defense resources.

Existing cybersecurity literature related to transportation infrastructure is limited, especially from the perspective of emergent behaviors and system level performance. This project explores the impact of cyber-attacks on transportation mobility and safety and coalesces knowledge from diverse fields such as optimization, machine learning, and statistical inference. The project aims to bridge research gaps between the cybersecurity and transportation disciplines by incorporating domain knowledge (e.g., traffic flow models) into security analysis.

Project URL: https://www.nsf.gov/awardsearch/showAward?AWD_ID=2339753&HistoricalAwards=false

# Collaborative Research: SaTC: Medium: Exposing and Mitigating Security/Safety Concerns of CAVs: A Holistic and Realistic Security Testing Platform for Emerging CAVs

**PIs: Yiheng Feng, Z. Morley Mao (University of Michigan), Q. Alfred Chen (University of California, Irvine)**

### Overview

Connected and Automated Vehicle (CAV) technologies enable real-time information sharing and driving automation, with the potential of significantly improving safety and efficiency of the transportation system. However, cyber-security threats may compromise the efficiency of infrastructure operations and the safety of passengers, posing a significant challenge for CAV deployment. This collaborative project develops a novel CAV testing platform to address the critical needs for assessing the security and safety concerns of the CAV system in an effective and realistic manner. Hardware manufacturers, software developers, and security service providers in the CAV industry can all leverage such a platform to conveniently and holistically test their products. Moreover, the testing platform can be used for training and education in both academia and industry, and facilitate the development of security best practices and standards in industry.

The testing platform provides both offensive and defensive testing services covering three parts of the CAV ecosystem: (1) transportation infrastructure, (2) connected vehicle communication channels, and (3) in-vehicle software platform. It is the first CAV security testing platform supporting cross-component offensive and defensive testing services designed to capture the inter-relationships among the three key components. The project develops testing support by effectively combining techniques in optimization, statistical modeling, machine learning, network emulation, program analysis, and model checking, while supporting innovative use of known defense solutions, including trusted computing hardware, access control policy enforcement, and anomaly detection. The project implements and deploys the testing platform in real-world CAV testbeds and collaborates with industry partners for early adoption.

Project URL: https://www.nsf.gov/awardsearch/showAward?AWD_ID=1930041

# Cloud-connected Cyber-physical Edge Devices for Electric Vehicles Eco-systems with Low Power and Cost

**PIs: Athula Kulatunga, Dr. Daniel Sampaio, UNESP, Sao Paqulo, Brazil, Dr. Uditha S. Navaratne, University of Peradeniya, Sri Lanka, Dr. Ajith Wijenayake, Senior Manager, HV Electronics, E-Motors and Charging Sys, Electrified Powertrain Propulsion Systems, Fiat-Chrysler Automobiles**

### Overview

Value-added services for individual devices in the BEV eco-system are possible when cloud-based services are used. Fault detection of electric drive motors to the power needed by a specific charging station in the charging infrastructure brings many operational and economic benefits. Building application-specific IoT platforms from the ground up, with the ability to integrated fault detection algorithms and machine learning capabilities in the edge devices, reduce the cost. Commercial cloud services can

cost prohibitive when frequent access and duplication increase. The current research builds upon the lessons learned through developing hardware infrastructure for Advanced Meter Infrastructure (AMI) security key management research in Smart Grids. The main focus is on building robust, secured, smart decision making physical infrastructure that subscribe/publish via low-cost channels to custom build IoT platforms. The cyber-physical system design encompasses improving collective system efficiency as well. Other possibilities include developing cyber-physical systems capable of meeting the maturity levels of various Industry 4.0 standards. A symposium was held in fall 2019 for gap analysis of Industry 4.0 and IoT platforms.

## Bootloader Security

**PI: Aravind Machiry**

*Funding Source: DARPA*

**Overview**
We are working on various (static + dynamic) techniques to secure bootloaders, especially UEFI/EDK-2.

## Securing Deeply Embedded Devices

**PIs: Aravind Machiry, Jamie Davis**

*Funding Source: National Science Foundation (NSF), Rolls Royce*

**Overview**
we are investigating ways to effectively test deeply embedded devices.

Project URL: https://machiry.github.io/research/

# Agricultural Robotic Systems for Greenhouses

**PI: Shimon Nof Current Students: Maitreya Sreeram, Puwadol (Oak) Dusadeerungsikul, Win Nguyen**
*Funding Source: Agricultural Research Organization, Volcani Center, Israel*

### Overview

Food security is an important challenge for society, especially in the face of climate change and increasing world population. The food security challenge is addressed by this project, which aims to develop the collaborative design and control principles and methodologies for agricultural robotic systems (ARS) for greenhouses. The ARS is designed to monitor, detect, and respond to different types of plant stresses to ensure reliable quality and productivity. This cyber-physical system involves various types of agents: human farmer experts, agricultural autonomous robots, and environmental sensors. These agents must collaborate in order to overcome the unstructured and everchanging environment of agricultural products.

### Representative Publications

Dusadeerungsikul, Puwadol Oak, Shimon Y. Nof, Avital Bechar, and Yang Tao. "Collaborative Control Protocol for Agricultural Cyber-Physical System." Procedia Manufacturing, ICPR-25, Chicago, IL August 2019.

Dusadeerungsikul, Puwadol Oak, and Shimon Y. Nof. "A collaborative control protocol for agricultural robot routing with online adaptation." Computers & Industrial Engineering (2019).

Guo, Ping, Puwadol Oak Dusadeerungsikul, and Shimon Y. Nof. "Agricultural cyber physical system collaboration for greenhouse stress management." Computers and electronics in agriculture 150 (2018): 439-454.

Project URL: https://engineering.purdue.edu/~prism/prj_ars.shtml

# ONR-BAA: Reactor Simulation Tool for Investigating the Resilience of a Cyberphysical Security Ecosystem

**PI: Lefteri Tsoukalas, Miltos Alamaniotis (Co-PI)**

**Current Students: Styliani Pantopoulou, Clive H. Townsend, Pola Lydia Lagari**

*Funding Source: Office of Naval Research (ONR)*

### Overview

Digitalization of critical systems has attracted a lot of research interest during the past years, as it offers numerous advantages, such as reduction of system complexity, ability of remote monitoring, reduction of costs etc. However, the incorporation of digital control to critical systems projects major concerns, the most challenging of which is that of vulnerability towards cyber-attacks. This work investigates how the physics-based nature of a nuclear reactor can be exploited in order to enhance the resilience of this system.

The main approach is to create a connection between the cyber-space actions with physical operational patterns in the system. This strategy can enable the early detection of cyber-attacks by monitoring the state of the system and by mapping cyber-events to physical faults, thus rendering the system resilient. A nuclear reactor model is built, with the purpose of simulating the actual system and monitoring its behavior towards targeted cyber-attacks. This simulation is a result of combining advanced reactor codes, such as SCALE/ORIGEN, MCNP, RELAP5 etc. with a control architecture. The final model is comprised of the physical and the control variables. Additionally, a Graphical User Interface (GUI) has a purpose of facilitating the interaction between the user and the system. The Purdue University Reactor-1 (PUR-1) serves as a testbed in order to perform benchmarking for model validation.

# End System Security

## Protecting and Securing Supply Chain Data throughout its Lifecycle

**PI: Bharat Bhargava**

### Overview

Enterprises operate in a global economy and their manufacturing and supply operations are dispersed throughout the world. This makes supply chain a critical and integral part and offers research challenges. There are multiple stages in supply chain and each stage generates data that is shared and transferred among different steps, divisions, or processes. In large enterprise systems, it is difficult to understand and track the sharing and dissemination of sensitive information. The sharing of data across multiple processes and divisions in a supply chain complicates and magnifies the problem further. The effect of shared data being compromised is one of the key risks in the supply chain. Vulnerabilities of business processes are as important as vulnerabilities in network attacks, viruses and malwares). Common existing approaches, standards, and guidelines ensure security but the focus is on the protection of data inside the private domain of an organization and do not address the protection in a decentralized supply chain. Current approaches rely on service level agreements or contracts, or perform audits to provide a level of assurance at the top level of their supply chain. In this project, we propose an end-to-end security framework that would ensure data protection (security and privacy) throughout its life cycle and enable organizations to securely share and disseminate data in their supply chains steps.

Project URL: http://www.cs.purdue.edu/homes/rranchal/plm.html

## NARCISSUS: Deriving Correct-By-Construction Decoders and Encoders from Binary Formats

**PI: Benjamin Delaware**
**Current Students: Qianchuan Ye**

### Overview

NARCISSUS is a framework for automatically synthesizing high-assurance serializer and deserializers from high-level specifications. The starting point of the process is a binary format, expressed as a relation which precisely captures all the valid binary encodings of an arbitrary datatype instance. From this specification, NARCISSUS synthesizes a decoder that is guaranteed to be the inverse of this relation, drawing on an extensible set of decoding strategies to construct the implementation. Each decoder is furthermore guaranteed to detect malformed encodings by failing on inputs not included in this relation. The derivation is carried out inside the Coq proof assistant and produces a proof trail certifying the correctness of the synthesized decoder. We have demonstrated the utility of our framework by deriving and evaluating the performance of decoders for all packet formats used in a standard network stack.

### Representative Publications

Benjamin Delaware, Sorawit Suriyakarn, Clément Pit-Claudel, Qianchuan Ye, and Adam Chlipala. 2019. Narcissus: correct-by-construction derivation of decoders and encoders from binary formats. Proc. ACM Program. Lang. 3, ICFP, Article 82 (August 2019), 29 pages. https://doi.org/10.1145/3341686

Qianchuan Ye and Benjamin Delaware. 2019. A verified protocol buffer compiler. In Proceedings of the 8th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP 2019). Association for Computing Machinery, New York, NY, USA, 222–233. https://doi.org/10.1145/3293880.3294105

## MicroVM: Micro Virtual Machines for Managed Languages – Abstraction, contained

**PI: Antony Hosking, Eliot Moss, UMass Amherst, Steve Blackburn, Australian National Univ., Michael Norrish, Data61**

*Funding Source: Australian Research Council, Data61, National Science Foundation (NSF)*

### Overview

A large fraction of today's software is written in managed languages. These languages increase software productivity by offering rich abstractions for managing memory, executing code concurrently, and hiding the complexity of modern hardware. Examples include JavaScript, PHP, Objective-C, Java, C#, Python, and Ruby. These languages are economically important. Unfortunately, most of these languages are inefficient, imposing overheads as large as a factor of fifty compared to orthodox language choices such as C.

The project will define, develop, evaluate, and refine the essential components of a new foundation layer for managed language implementation. In doing so, it will address a key source of systemic inefficiency, by pioneering the micro virtual machine (μVM) as an efficient high-performance substrate for managed language implementation. The relationship between a μVM and existing managed language implementations is analogous to the one between an operating system micro kernel and monolithic operating systems such as Linux. A μVM captures the insight that there exists a well-defined foundation common to most modern languages that can take responsibility for fundamental abstractions over hardware, concurrency, and memory. By isolating and exposing this substrate, a μVM embodies state-of-the-art base technology available to language implementers while isolating them from the pernicious complexities of these abstractions, freeing them to focus on all-important language-specific optimizations. This project will enable more efficient, reliable, and verifiable software, and a distinctly sharper focus for language implementation research and development.

### Representative Publications

Y. Lin, S. M. Blackburn, A. L. Hosking, and M. Norrish. Rust as a language for high performance GC implementation. In ACM SIGPLAN International Symposium on Memory Management, ISMM, pages 89–98, Santa Barbara, California, June 2016. doi: 10.1145/2926697.2926707

P. Gammie, A. L. Hosking, and K. Engelhardt. Relaxing safely: Verified on-the-fly garbage collection for x86-TSO. In ACM SIGPLAN International Conference on Programming Language Design and Implementation, PLDI, pages 99–109, Portland, Oregon, June 2015b. doi: 10.1145/2737924.2738006

Y. Lin, K. Wang, S. M. Blackburn, A. L. Hosking, and M. Norrish. Stop and go: Understanding yieldpoint behavior. In ACM SIGPLAN International Symposium on Memory Management, ISMM, pages 70–80,

Portland, Oregon, June 2015. doi: 10.1145/2754169.2754187

K. Wang, Y. Lin, S. M. Blackburn, M. Norrish, and A. L. Hosking. Draining the swamp: Micro virtual machines as solid foundation for language development. In T. Ball, R. Bodík, S. Krishnamurthi, B. S. Lerner, and G. Morrisett, editors, Inaugural Summit on Advances in Programming Languages, SNAPL, pages 321–336, Asilomar, California, May 2015. doi: 10.4230/LIPIcs.SNAPL.2015.321
Project URL: http://microvm.org

## PeX: A Permission Check Analysis Framework for Linux Kernel

**PI: Changhee Jung**

### Overview

Permission checks play an essential role in operating system security by providing access control to privileged functionalities. However, it is particularly challenging for kernel developers to correctly apply new permission checks and to scalably verify the soundness of existing checks due to the large codebase and complexity of the kernel. In fact, Linux kernel contains millions of lines of code with hundreds of permission checks, and even worse its complexity is fast-growing.

This work presents PeX, a static Permission check error detector for LinuX, which takes as input a kernel source code and reports any missing, inconsistent, and redundant permission checks. PeX uses KIRIN (Kernel InteRface based In-direct call aNalysis), a novel, precise, and scalable indirect call analysis technique, leveraging the common programming paradigm used in kernel abstraction interfaces. Over the inter-procedural control flow graph built by KIRIN, PeX automatically identifies all permission checks and infers the mappings between permission checks and privileged functions. For each privileged function, PeX examines all possible paths to the function to check if necessary permission checks are correctly enforced before it is called.

We evaluated PeX on the latest stable Linux kernel v4.18.5 for three types of permission checks: Discretionary AccessControls (DAC), Capabilities, and Linux Security Modules(LSM). PeX reported 36 new permission check errors, 14 of which have been confirmed by the kernel developers.

### Representative Publications

Tong Zhang, Wenbo Shen, Dongyoon Lee, Changhee Jung, Ahmed Azab, and Ruowen Wang
“PeX: A Permission Check Analysis Framework for Linux Kernel”,
The 28th USENIX Security Symposium (USENIX Security), Santa Clara, California, August 2019.

## Rust for Embedded Systems

*PI: Aravind Machiry*

### Overview

We are working on techniques to enable using Rust for embedded systems.
https://machiry.github.io/files/rustembedccs.pdf

# Assurable Configuration of Security Policies in Enterprise Networks

**PI: Sanjay Rao** **Current Students: Xin Sun, Yu-Wei Sung**

### Overview

The design and configuration of enterprise networks is one of the hardest challenges that operators face today. A key challenge in doing so is the need to reconfigure network devices to ensure high-level operator goals are correctly realized. The high-level objectives (such as performance and security goals) that operators have for their networks are embedded in hundreds of low-level device configurations. Reconfiguring network devices is challenging given the huge semantic gap between these high-level objectives, and low-level configurations. Errors in changing configurations have been known to result in outages, business service disruptions, violations of Service Level Agreements~(SLA) and cyber-attacks~\ cite{mahajan:02,kerravala02,Alloy}. In our research, we are looking at principled approaches for the systematic design and configuration of enterprise networks. We believe our research will minimize errors, and enable operators to ensure their networks continue to meet desired high-level security objectives. An important problem that we are currently tackling is that of ensuring correctness of security policies when migrating enterprise data centers to cloud computing models.

### Representative Publications

Towards Systematic Design of Enterprise Networks, Yu-Wei Eric Sung, Sanjay Rao, Geoffrey Xie, and David Maltz. Proceedings of ACM CoNEXT, Madrid, Spain, December, 2008.

Modeling and Understanding End-to-end Class of Service Policies in Operational Networks, Yu-Wei Sung, Carsten Lund, Mark Lyn, Sanjay Rao, Shubho Sen, Proceedings of ACM SIGCOMM, Barcelona,Spain, August 2009

Project URL: http://www.ece.purdue.edu/~isl

# Migrating Enterprises to Hybrid Cloud Architectures

**PI: Sanjay Rao**

**Current Students: Mohammad Hajjat, Xin Sun, Yu-Wei Sung**

### Overview

We are tackling challenges in migrating enterprise services into hybrid cloud-based deployments, where enterprise operations are partly hosted on-premise and partly in the cloud. Such hybrid models enable enterprises to benefit from cloud-based architectures, while honoring privacy and other restrictions on what data and services may be migrated to the cloud. We are making several contributions. First, we are obtaining deeper understanding about the architecture of enterprise applications today in terms of their multi-tiered nature, large number of application components, and interdependencies, and security policies associated with enterprise applications in data centers. Second, we are showing the need for and benefits of a planned approach to deciding what application components to migrate to the cloud. Our approach is based on framing and solving an optimization problem that takes into End System Security account enterprise-specific constraints, cost savings from migration, and increased transaction delays and wide-

area communication that may result from the migration. Third, we are developing algorithms to ensure security policies are correctly reconfigured as enterprise applications are migrated to the cloud. Our evaluations are conducted using models of real enterprise applications and data center security policies derived from discussions with operators in a university setting, and using actual Azure-based cloud deployments. The results are promising and show the potential of our approach.

## Cryptanalysis of RSA

**PIs: Samuel Wagstaff, Peter L. Montgomery**

**Current Students: Sangil Nahm**

**Overview**

*Funding Source: Microsoft*

We study the minimum period of the Bell numbers, which arise in combinatorics, modulo a prime. It is shown that this period is probably always equal to its maximum possible value. Interesting new divisibility theorems are proved for possible prime divisors of the maximum possible period. The conclusion is that these numbers are not suitable for use as RSA public keys.

*Representative Publications*

Peter L Montgomery, Sangil Nahm and Samuel S Wagstaff Jr., "The period of the Bell numbers modulo a prime," Math. Comp. v. 79 (2010), pp. 1793--1800.

Project URL: http://homes.cerias.purdue.edu/~ssw/bell/index.html

## Search for Aurifeuillian Factorizations

**PI: Samuel Wagstaff, Mikhail Atallah**

**Current Students: Paul Kuliniewicz, Usman Latif**

**Overview**

We searched the online Cunningham tables for new algebraic factorizations similar to those discovered by Aurifeuille. A naive search would have taken too long. We accelerated it enough to make the search feasible. Many interesting results were found.

Project URL: http://homes.cerias.purdue.edu/~ssw/cun/index.html

# Knowledge Graph Construction for Resilient, Trustworthy, and Secure Software Supply Chains

**PIs: Tianyi Zhang, Xiangyu Zhang**

**Current Students: Yifeng Di, Yuan Tian, Bonan Kou, Minghai Lu, Zhi Tu, Ruixin Wang, Weixi Tong, Jiahao Shi, Wei-Hao Chen**

## Overview

This project will develop a unified knowledge graph that captures rich, up-to-date information about software components in heterogeneous software ecosystems. Building upon our prior work on noise-robust open knowledge extraction, we will develop a new neural knowledge acquisition pipeline that (1) extracts software information from various information sources, including but not limited to official documentation, software release notes, bug reports, CVEs, and online discussions, (2) consolidates the extracted information via an array of quality control and fact-checking mechanisms, and (3) constantly updates the knowledge graph by tracking new information from various sources. The resulting knowledge graph will empower us to further develop a novel multi-modal query interface for knowledge dissemination, as well as new risk mitigation approaches that perform deep scans on software systems, detect potential risks, and automatically repair them.

## Representative Publications

Zhao, Zhongkai, Bonan Kou, Mohamed Yilmaz Ibrahim, Muhao Chen, and Tianyi Zhang. "Knowledge-Based Version Incompatibility Detection for Deep Learning." *ESEC/FSE 2023*.

Nguyen, Tai, Yifeng Di, Joohan Lee, Muhao Chen, and Tianyi Zhang. "Software Entity Recognition with Noise-Robust Learning." *ASE 2023*.

Tian, Yuan, Zheng Zhang, Zheng Ning, Toby Li, Jonathan K. Kummerfeld, and Tianyi Zhang. "Interactive Text-to-SQL Generation via Editable Step-by-Step Explanations." In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 16149-16166. 2023.

Yuan Tian, Jonathan K. Kummerfeld, Toby Jia-Jun Li, and Tianyi Zhang. 2024. SQLucid: Grounding Natural Language Database Queries with Interactive Explanations. In Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology (UIST '24). Association for Computing Machinery, New York, NY, USA, Article 12, 1–20. https://doi.org/10.1145/3654777.3676368

Zhou, Zihan, Zhongkai Zhao, Bonan Kou, and Tianyi Zhang. "Decide: Knowledge-Based Version Incompatibility Detection in Deep Learning Stacks." In *Companion Proceedings of the 32nd ACM International Conference on the Foundations of Software Engineering*, pp. 547-551. 2024.

Project URL: https://purdue-hcss.github.io/nsf-software-supply-chain_security/

# Convicting Exploitable Software Vulnerabilities: Practical Input Provenance-Based Approach

**PIs: Xiangyu Zhang, Dongyan Xu**

## Overview

*Funding Source: National Science Foundation (NSF)*

Vulnerabilities in software, especially those that are remote exploitable, are the root cause of wave after wave of security attacks, such as botnet, zero-day worms, non-control data corruptions, and even server-break-ins. Thus, analyzing and exposing software vulnerabilities has become one of the most active research areas today. In the past, software vulnerability detection/exposing approaches could be divided into two categories: dynamic and static. Static analysis creates a lot of false positives. Dynamic approaches monitor program execution and detect attempts of attacking a software system. These techniques incur non-trivial runtime overhead and cannot detect vulnerabilities that not under attack. Dynamic test generation has the potential of generating exploit inputs to confirm vulnerabilities. However, most existing dynamic test generation techniques suffer from the scalability problem. In this project, we develop a practical dynamic approach that is intended to use in combination with other static tools. We observe that although the suspect pool produced by existing static tools has a high false positive rate, it is nonetheless much smaller than the whole population. Therefore, we use existing static tools as the frontend to generate a set of suspects. Our technique then tries to generate exploits for these suspects. A suspect is convicted only when an exploit can be acquired as the evidence. Such exploits significantly assist regular users and administrators to evaluate the robustness of their software and convince vendors to debug and patch.
The key idea is to use data lineage tracing to identify a set of input values relevant to the execution of a vulnerable code location. Exploit specific mutations are applied to the relevant input values in order to trigger an attack, e.g., for example, changing an integer value to MAXUINT to induce an integer overflow. Since these inputs are usually a very small subset of the whole input sequence, mutating the whole input, like in random test generation, is avoided. Our technique does not rely on symbolic execution and constraint solving and thus can easily handle long execution. In case an execution that covers a vulnerable code location cannot be found, our technique also allows user interactions to mutate an input so that the execution driven by the mutated input covers the vulnerable code location. Our technique addresses a wide range of vulnerabilities including buffer overflow, integer overflow, format string, etc. Our dynamic analysis works at binary level, which greatly facilitates users that do not have the source code access but are concerned about software vulnerabilities. We have developed a data lineage tracing prototype. It traces the set of input that is relevant to a particular execution point. The lineage
information is used to guide our evidence generation procedure. The challenge of efficiency is overcome by using Reduced Ordered Binary Decision Diagrams (RoBDDs). Our initial experience with a set of known and unknown real vulnerabilities showed that our technique can very quickly generate exploit inputs.

## Representative Publications

Z. Lin, X. Zhang, and D. Xu, Convicting Remote Exploitable Vulnerabilities: An Efficient Input Provenance Based Approach, Proceedings of IEEE/IFIP International Conference on Dependable Systems and Networks (DSN), 2008.

Project URL: http://http: //www.cs.purdue.edu/homes/xyzhang/Comp/dsn08.pdf

# Secure Group Communication Over Wired/Wireless Networks

**PIs: Xukai Zou, Byrav Ramamurthy (UNL), V. N. Variyam (UNL)**

## Overview

*Funding Source: National Science Foundation (NSF)*

Secure group communications (SGC) refers to a setting in which a group of participants can send and receive messages (sent to the group members), in a way that outsiders are unable to glean information even if they are able to intercept the messages. SGC is important because several prevalent applications require it. These applications include teleconferencing, tele-medicine, real-time information services, distributed interactive simulations, collaborative work, interactive games and the deployment of VPN (Virtual Private Networks). The goals for this project are four-fold:

- Study various issues enabling SGC which include, but are not limited to, group key management, burst behavior and efficient burst operations, group membership management, group member admission control, authentication and non-repudiation;
- Study and provide solutions for specific SGC scenarios such as dynamic conferencing and SGC with hierarchical access control;
- Investigate research challenges for SGC over wireless/mobile environments;
- Integrate research results into the curriculum and perform public dissemination of findings and software.

## Representative Publications

X. Zou, B. Ramamurthy and Spyros Magliveras, "Secure Group Communication over Data Networks", (Oct. 2004). Springer, ISBN: 0-387-22970-1.

P. Adusumilli, X. Zou and B. Ramamurthy, DGKD: Distributed Group Key Distribution with Authentication Capability, Proceedings of the 2005 IEEE Workshop on Information Assurance (IAW), United States Military Academy, West Point, NY, 15-17 June 2005, pp. 286–293.

X. Zou, A. Thukral, and B. Ramamurthy, An Authenticated Key Agreement Protocol for Mobile Ad Hoc Networks, LNCS, Springer, Vol. 4325, pp. 509–520, 2006.

Y.Wang, B. Ramamurthy, and X. Zou, KeyRev: An Efficient Key Revocation Scheme for Wireless Sensor Networks, Proceedings of IEEE ICC'07, 24-27 June 2007, Scotland, pp.1260 - 1265.

X. Zou, Y. Dai and E. Bertino, A Practical and Flexible Key Management Mechanism For Trusted Collaborative Computing, Proceedings of the 27th IEEE INFOCOM, April 13–18, 2008, pp. 1211–1219.

R. Balachandran, X. Zou, B. Ramamurthy, A. Thurkral, and N. V. Vinodchandran, An Efficient and Attack-resistant Key Agreement Scheme for Secure Group Communications in Mobile Ad-Hoc Networks, Wireless Communications & Mobile Computing. 8(10), 2008, pp. 1297-1312.

X. Zou, Y. Dai and Y. Pan, Trust and Security in Collaborative Computing, World Scientific, ISBN-13: 978-981-270-3682, January 2008.

# Human Centric Security

## Securing IoT-based Cyber-Physical Human Systems against Collaborative Attacks

**PIs: Bharat Bhargava, Sathish A.P Kumar at Coastal Carolina University, Conway, SC, USA, Raimundo Macêdo at Federal University of Bahia, Ondina, Salvador, Bahia, Brazil**

**Current Students: Ganapathy Mani**

### Overview

Security issues in the IoT-based Cyber-Physical Systems (CPS) are exacerbated with human participation in CPHS due to the vulnerabilities in both the technologies and the human involvement. A holistic framework to mitigate security threats in the IoT-based Cyber-Physical Human Systems (CPHS) environment is presented to mitigate these issues. We have developed threat model involving human elements in the CPHS environment. Research questions, directions, and ideas with respect to securing IoT-based CPHS against collaborative attacks are presented.

Project URL: https://www.cs.purdue.edu/homes/bb/#colloquia

## RUDOLF: An Efficient and Adaptive Defense Approach Against Website Fingerprinting Attacks Based on Soft Actor-Critic Algorithm

**PI: Bharat Bhargava, Junsong Fu**

**Current Students: Meiyi Jiang; Baojiang Cui; Tao Wang; Lu Yao; Bharat K. Bhargava, All Authors**

### Overview

Although Tor is designed to provide anonymity, website fingerprinting (WF) attacks have posed significant threats to user privacy. In response, various defense approaches have been developed. Randomization and regularization-based defenses are criticized to be inefficient due to their bandwidth-consuming nature. Some adversarial learning-based defenses are impractical because the generation of perturbation depends on the complete traffic traces. Other adversarial learning-based defenses have weaknesses of lacking adaptability because their perturbations are input-agnostic. To overcome these shortcomings, we propose RUDOLF, an efficient and adaptive WF defense based on the soft actor-critic (SAC) algorithm of reinforcement learning (RL). We train the agent that can incrementally output perturbations synchronously following each burst of real-time traffic. Different from previous defenses, RUDOLF's perturbation does not depend on the integrity of the traffic and concerns the actual real-time traffic, which ensures the practicality of implementation and adaptability. Besides, we take advantage of the exploratory characteristics of the SAC algorithm to obtain the optimal policy of adding perturbations that can efficiently balance defense effects and bandwidth consumption. Experiments on synthetic datasets show that with less than 30% bandwidth overhead (BWO), RUDOLF can reduce the average attack accuracy to around 15%–20%, which is superior to previous works. We also have implemented RUDOLF as a Tor

pluggable transport. The performance in the real Tor network shows that RUDOLF can reduce the average accuracy of WF classifier to around 24% with about 25% BWO and almost no time delay.

### Representative Publications

M. Jiang, B. Cui, J. Fu, T. Wang, L. Yao and B. K. Bhargava, "RUDOLF: An Efficient and Adaptive Defense Approach Against Website Fingerprinting Attacks Based on Soft Actor-Critic Algorithm," in IEEE Transactions on Information Forensics and Security, vol. 19, pp. 7794-7809, 2024, doi: 10.1109/TIFS.2024.3436818

Project URL: https://ieeexplore.ieee.org/document/10620322

## Economics of Password Cracking

**PI: Jeremiah Blocki**

**Current Students: Ben Harsha, Samson Zhou**

### Overview

In recent years the authentication servers at major companies like Yahoo!, Dropbox, Ashley Madison, eBay, Zappos, Sony, LinkedIn and Adobe have been breached resulting in the release of the cryptographic hashes of over a billion of user passwords, each of which has significant economic value to adversaries. An adversary who has obtained the cryptographic hash of a user's password can mount an offline attack to crack the password by comparing this hash value with the cryptographic hashes of likely password guesses. This offline attacker is limited only by the resources he is willing to invest to crack the password. This project aims to address the following questions: Can we quantitatively predict how many passwords a rational attacker will crack after a breach? What defenses could an authentication server adopt to minimize the potential damage if a password breach does occur?

We have developed an economic model of password cracking, which are currently using to analyze recent password breaches. (e.g., what % of Yahoo! will be cracked by an attacker?). The basic premise of our model is that a rational attacker should cease attacking once marginal guessing costs exceed the marginal guessing reward. An attacker's marginal guessing reward is parameterized by the value of a cracked password and by the empirical distribution over user selected passwords (e.g., expected marginal reward is the value a cracked password times the probability that the next password guess is correct). Marginal guessing costs are given by the (amortized) cost of computing the password hash function. Most password hashing algorithms (e.g., Argon2i, SCRYPT, BCRYPT) have time/memory parameters which can be tuned to increase/decrease marginal guessing costs. We are currently working on using our economic model to provide concrete recomendations about how to set the parameters of a password hashing algorithm to ensure that the % of passwords which would be cracked by an offline attacker in the event of a breach is acceptably small.

### Representative Publications

CASH: A Cost Asymmetric Secure Hash Algorithm for Optimal Password Protection. with Anupam Datta. CSF 2016.

Client-CASH: Protecting Master Passwords against Offline Attacks. with Anirudh Sridhar. AsiaCCS 2016.

On the Economics of Offline Password Cracking. with Ben Harsha and Samson Zhou. (Working Paper)

## Security and Privacy in Augmented and Virtual Reality (AR/VR)

**PI: Berkay Celik**

**Overview**
This project investigates emerging threats in mixed reality environments. The research includes analyzing user interface (UI) attacks in WebXR, exploring the extraction of sensitive information like speech from zero-permission VR sensors, and inferring users' semantic locations from the spatial maps generated by their devices. The primary objective is to identify and create defenses against vulnerabilities that are unique to AR/VR platforms.

Project URL: https://beerkay.github.io/

## A Human Factors Perspective on Better Phishing Defenses

**PI: Jamie Davis**

**Current Students: Andrew Rozema, PhD student (he is also a professor)**

**Overview**
Social engineering attacks delivered via email, commonly known as phishing, represent a persistent cybersecurity threat leading to significant organizational incidents and data breaches. Although many organizations train employees on phishing, often mandated by compliance requirements, the real-world effectiveness of this training remains debated. Past work has demonstrated the ineffectiveness of training, but reproduction across different organizations, training approaches, and with a standardized threat assessment will help the generalizability of this phenomenon.

This project has several goals:

6 Measure the effectiveness of state-of-art phishing trainings in real-world settings
7 Explore opportunities for email defense system optimization through improved and faster feedback loops
8 Trial new phishing attacks in real-world settings
9 Representative Publications

Anti-Phishing Training (Still) Does Not Work: A Reproduction of Phishing Training Inefficacy Grounded in the NIST Phish Scale, Rozema & Davis, arXiv'25

# High-fidelity and Trustworthy Teleinteraction Platform

**PI: Younghyun Kim**

*Funding Source: Ministry of Science and ICT, South Korea*

**Overview**
This research aims to develop a cutting-edge teleinteraction platform that enables highly immersive, realistic, and trustworthy virtual experiences by integrating three core technologies: (1) leveraging multimodal sensors and AI models to analyze and virtualize high-precision mixed reality spaces, (2) establishing a real-time, low-power interaction system using devices such as head-mounted displays and wearables, and (3) developing authentication and security technologies for real-time detection and verification of mixed reality spaces and user information.

# Understanding the Impacts of Human Decision-Making on Security and Robustness of Large-Scale Systems

**PIs: Shreyas Sundaram, Saurabh Bagchi, Timothy Cason**

**Current Students: Ashish Ranjan Hota, Mustafa Abdallah El-Hosiny**

*Funding Source: Internal, National Science Foundation (NSF)*

**Overview**
The robustness and security of systems depend critically on the way they are utilized by human decision-makers. While there are various classical mathematical frameworks that have been used to model decision-makers, studies in behavioral psychology and economics have shown that humans consistently deviate from such traditional models of behavior. These deviations from expected behavior, particularly in the way that humans view and evaluate risks and losses, can significantly impact the way that they use shared systems.

In this project, we consider the game-theoretic implications of behavioral deviations (captured by Prospect Theory) on the utilizations of failure- and attack-prone systems. For example, we consider a setting where a group of individuals utilize a shared resource; the resource fails with a probability that increases with the amount of utilization, and provides a certain return otherwise. We show that utilization increases as the players deviate from risk neutrality, and also when they have heterogeneous attitudes towards loss. We also consider the use of taxation policies to mitigate overutilization of the resource, and demonstrate that counter-intuitive outcomes can arise under behavioral decision-making.

We also consider interdependent security games where each node in a network chooses how much to invest in security to protect itself. The successful attack probability at each node in such settings depends on the investment at that node and on neighboring nodes. We characterize the impact of prospect-theoretic perceptions of attack probabilities on the equilibrium security investments, and identify techniques to optimally design networks to mitigate security risks under behavioral decision-making.

### Representative Publications

A. R. Hota, S. Garg and S. Sundaram, "Fragility of the Commons under Prospect-Theoretic Risk Attitudes." Games and Economic Behavior, vol. 98, pp. 135 - 164, July 2016.

A. R. Hota and S. Sundaram, "Optimal Network Topologies for Mitigating Security and Epidemic Risks." Proceedings of the 54th Annual Allerton Conference on Communication, Control, and Computing, Allerton, IL, 2016 (invited).

A. R. Hota and S. Sundaram, "Controlling Human Utilization of Shared Resources via Taxes." Proceedings of the 55th IEEE Conference on Decision and Control, Las Vegas, NV, 2016.

A. R. Hota and S. Sundaram, "Interdependent Security Games under Behavioral Probability Weighting." Proceedings of GameSec 2015, the Conference on Decision and Game Theory for Security, London, England, 2015.

A. R. Hota, A. A Clements, S. Sundaram and S. Bagchi, "Optimal and Game-Theoretic Deployment of Security Investments in Interdependent Assets." Proceedings of GameSec 2016, the Conference on Decision and Game Theory for Security, New York City, NY, 2016.

A. R. Hota and S. Sundaram, "Game-Theoretic Protection Against Networked SIS Epidemics by Human Decision-Makers." Proceedings of the 2nd IFAC Workshop on Cyber-Physical and Human Systems, Miami, FL, December 2018 (to appear).

Project URL: https://engineering.purdue.edu/~sundara2/index.html

## Secure, Composable, & Scalable Framework for Trusted Collaborative Computing

**PI: Xukai Zou**

### Overview

Collaborative Computing (CC) is a critical application domain within the Internet environment. A few examples of CC are multi-party computation, collaborative defense, tele-medicine and collaborative decision making. Participants in CC demand confidentiality, privacy, integrity, and controlled sharing of sensitive information. Also, CC environments involve many entities, which are dynamic, heterogeneous, distributed, and can be hostile. Currently, CC uses the Internet as the underlying infrastructure, which by design is not secure and suffers from incessant attacks ranging from eavesdropping to vulnerability exploitation. Hence, it is imperative for the success of CC to require a reliable and secure framework built on top of the Internet to remedy some of its limitations. CC, based on such an underlying framework, can be termed as Trusted Collaborative Computing (TCC). Thus, the long term objective of this research is to develop a framework that will enable TCC. This framework consists of: (1) (group-oriented) secure and anonymous communication, (2) finely-controlled data sharing and (3) secure, composable and scalable integration. The framework will effectively address the underlying challenges of secure communication

and guaranteed access, anonymity, composability, interoperability, and scalability.

The core technique in the proposed TCC framework is Access Control Polynomial (ACP) which was just presented at and published in the proceedings of INFOCOM'08, one of the highest international conferences in the networking and security field. The short term yet intensive summer work is to implement and evaluate such an innovative ACP mechanism and related security modules. This work will significantly help the accomplishment of the long term objective and secure the application for external funding.

## Internet Based Electronic Voting Enabling Open and Fair Elections

**PI: Xukai Zou**

### Overview

Voting is the pillar of modern democracies. However, examination of current voting systems (including E-voting techniques) shows a gap between casting secret ballots and tallying and verifying individual votes. This gap is caused by either disconnection between the vote-casting process and the vote-tallying process, or opaque transition (e.g. due to encryption) from vote- casting to vote-tallying and thus, damages voter assurance, i.e. failing to answer the question: "Will your vote count?" We proposed a groundbreaking E-voting protocol that fills this gap and provides a fully transparent election. In this new voting system, this transition is seamless, viewable, and verifiable. As a result, the above question can be answered assuredly: "Yes, my vote counts!"

The new E-voting protocol is fundamentally different from all existing voting/E-voting pro- tocols in terms of both concepts and the underlying mechanisms. It consists of three innovative Technical Designs: TD1: universal verifiable voting vector; TD2: forward and backward mutual lock voting; and TD3: in-process verification and enforcement. The new technique is the first fully transparent E-voting protocol which fills the aforementioned gap. The trust is split equally among all tallying authorities who are of conflict-of-interest and will technologically restrain from each other. As a result, the new technique enables open and fair elections, even for minor or weak political parties. It is able to mitigate errors and risk and detect fraud and attacks including collusion, with convincingly high probability $1 - 2^{-(m-\log(m))}$ n (n: #voters and $m \geq 2$: #candidates). It removes many existing requirements such as trusted central tallying authorities, tailored hardware or software, and complex cryptographic primitives. In summary, the new e- voting technique delivers voter assurance and can transform the present voting booth based voting and election practice. Besides voting and elections, the new technique can also be adapted to other applications such as student class evaluation, rating and reputation systems.

# Revocable, Interoperable and User-Centric (Active) Authentication Across Cyberspace

PI: Xukai Zou

## Overview

This work addresses fundamental and challenging user authentication and universal identity issues and solves the problems of system usability, authentication data security, user privacy, irrevocability, interoperability, cross-matching attacks, and post-login authentication breaches associated with existing authentication systems. It developed a solid user-centric biometrics- based authentication model, called Bio-Capsule (BC), and implemented an (active) authentication system. BC is the template derived from the (secure) fusion of a user's biometrics and that of a Reference Subject (RS). RS is simply a physical object such as a doll or an artificial one, such as an image. It is users' BCs, rather than original biometric templates, that are utilized for user authentication and identification. The implemented (active) authentication system will facilitate and safely protect individuals' diffused cyber activities, which is particularly important nowadays, when people are immersed in cyberspace.

# Network Security

## Explainable AI Methods for Enhancing AI-Based Network Intrusion Detection Systems

**PI: Mustafa Abdallah**

**Current Students: Osvaldo Arreche**

**Overview**

In network security, the exponential growth of intrusions stimulates research toward developing advanced artificial intelligence (AI) techniques for intrusion detection systems (IDS). However, the reliance on AI for IDS presents challenges, including the performance variability of different AI models and the lack of explainability of their decisions, hindering the comprehension of outputs by human security analysts. Hence, this thesis proposes
end-to-end explainable AI (XAI) frameworks tailored to enhance the understandability and performance of AI models in this context.

The project first benchmarks seven black-box AI models across one real-world and two benchmark network intrusion datasets, laying the foundation for subsequent analyses. Subsequent work delves into feature selection methods, recognizing their crucial role in enhancing IDS performance by extracting the most significant features for identifying anomalies in network security. Leveraging XAI techniques, novel feature selection methods
are proposed, showcasing superior performance compared to traditional approaches. Also, this project introduces an in-depth evaluation framework for black-box XAI-IDS, encompassing global and local scopes. Six evaluation metrics are analyzed, including descriptive accuracy, sparsity, stability, efficiency, robustness, and completeness, providing insights into the limitations and strengths of current XAI methods. Finally, the project addresses the potential of ensemble learning techniques in improving AI-based network intrusion detection by proposing a two-level ensemble learning framework comprising base learners and ensemble methods trained on input datasets to generate evalua tion metrics and new datasets for subsequent analysis. Feature selection is integrated into both levels, leveraging XAI-based and Information Gain-based techniques.

Holistically, this project offers a comprehensive approach to enhancing network intrusion detection through the synergy of AI, XAI, and ensemble learning techniques by providing open-source codes and insights into model performances. Therefore, it contributes to the security advancement of interpretable AI models for network security, empowering security analysts to make informed decisions in safeguarding networked systems.
Representative Publications

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

XAI-IDS: Towards Proposing an Explainable Artificial Intelligence Framework for Enhancing Network Intrusion Detection Systems

# A Framework of High-Speed Network Protocol Fuzzing Based on Shared Memory

**PI: Bharat Bhargava, Junsong Fu**

**Current Students: Junsong Fu; Shuai Xiong; Na Wang; Ruiping Ren; Ang Zhou;**

### Overview

In recent years, security test of network protocols based on fuzzing has been attracting more and more attentions. This is very challenging compared with the stateless software fuzzing and most early network protocol fuzzers are of low speed and poor test effect. Since the first greybox and stateful fuzzer named AFLNET was proposed, several new schemes have been designed to improve its performance from different aspects. During the research, a great challenge is how to greatly improve the fuzzing efficiency. Based on the basic analysis in SNPSFuzzer, this article provides a more thorough analysis about the time consumption in a fuzzing iteration for 13 network protocols and then we design a High-speed Network Protocol Fuzzer named HNPFuzzer. In HNPFuzzer, the test cases and response messages between the client and server are transmitted through the shared memory, guided by a precise synchronizer, rather than the socket interfaces. This greatly shorten the period of an iteration. Moreover, we design a persistent mode attempting to fuzz the service instances in the memory more than one time based on analyzing the side effect information. This mode further improves the speed of fuzzing. Experiment results illustrate that our scheme can improve the fuzzing throughput by about 39.66 times in average and triggers a large number of crashes including 2 new vulnerabilities which cannot discovered by existing fuzzers. Note that, the existing network protocol fuzzing schemes proposed in different directions do not compete with each other and on the contrary, they can collaborate with each other to improve the overall fuzzing effect and efficiency. Consequently, more existing tools can be integrated into our framework to get better network protocol fuzzing effect.

### Representative Publications

J. Fu, S. Xiong, N. Wang, R. Ren, A. Zhou and B. K. Bhargava, "A Framework of High-Speed Network Protocol Fuzzing Based on Shared Memory," in IEEE Transactions on Dependable and Secure Computing, vol. 21, no. 4, pp. 2779-2798, July-Aug. 2024, doi: 10.1109/TDSC.2023.3318571
IEEE Transactions on Dependable and Secure Computing ( Volume: 21, Issue: 4, July-Aug. 2024)

Project URL: https://ieeexplore.ieee.org/document/10262045

# Adaptable Safety and Security in V2X Systems

**PI: Bharat Bhargava**

**Current Students: Miguel Villarreal-Vasquez, Pelin Angin**

### Overview

With the advances in the areas of mobile computing and wireless communications, V2X systems have become a promising technology enabling deployment of applications providing road safety, traffic efficiency and infotainment. Due to their increasing popularity, V2X networks have become a major target for attackers, making them vulnerable to security threats and network conditions, and thus affecting the

safety of passengers, vehicles and roads. Existing research in V2X does not effectively address the safety, security and performance limitation threats to connected vehicles, as a result of considering these aspects separately instead of jointly. In this work, we focus on the analysis of the tradeoffs between safety, security and performance of V2X systems and propose a dynamic adaptability approach considering all three aspects jointly based on application needs and context to achieve maximum safety on the roads using an Internet of vehicles. Experiments with a simple V2V highway scenario demonstrate that an adaptive safety/security approach is essential and V2X systems have great potential for providing low reaction times. Keywords-V2X systems; safety; security; authentication; adaptability.

Project URL: https://www.cs.purdue.edu/homes/bb/#colloquia

## BioKA-ASVN: Biometric-Based Key Agreement Scheme for Air Smart Vehicular Networks Using Blockchain Service

**PI: Bharat Bhargava, Ashok Kumar Das; David K. Y. Yau; Pascal Lorenz**

**Current Students: Basudeb Bera; Abhishek Bisht;**

### Overview

Air Smart Vehicular Networks (ASVNs) have become increasingly essentials in military and commercial industries in recent years, where drones are deployed to interact among each other via wireless medium in airspace due to their agility and versatility. ASVNs can form a closed loop from data perception to final execution by integrating communication devices, computation tools, and control modules. However, the used unencrypted and publicly available navigational signals like Global Positioning System (GPS) and communication over (public) insecure Internet connections, including wireless networks and WiFi, make ASVNs vulnerable to security breaches. Attackers can easily gain access to a drone's configuration and take control remotely, by launching various attacks such as GPS spoofing, false sensor data injection, maldrone malware injection, eavesdropping, man-in-the-middle, Denial-of-Service (DoS), replay, forgery, and Skyjack attacks. This paper proposes a robust and efficient multi-factor biometric-based security mechanism using blockchain as a service to address the security and privacy breaches in ASVNs. A comparative study demonstrates that the proposed scheme provides superior security and better functionality features with low communication and computation overheads as compared to the existing analogous schemes. The feasibility of the proposed scheme in real-life drone applications is also demonstrated through a real-time testbed and blockchain simulation. Furthermore, a detailed security analysis using the automated software validation tool, namely Scyther, verifies the proposed scheme's significant level of security.

### Representative Publications

B. Bera et al., "BioKA-ASVN: Biometric-Based Key Agreement Scheme for Air Smart Vehicular Networks Using Blockchain Service," in IEEE Transactions on Vehicular Technology, vol. 73, no. 7, pp. 9478-9494, July 2024, doi: 10.1109/TVT.2024.3380392.

Project URL: https://ieeexplore.ieee.org/document/10477615

# Quantitative Evaluation of Extensive Vulnerability Set Using Cost Benefit Analysis

**PIs: Bharat Bhargava, Urvashi Bansal; Geeta Sikka; Lalit K. Awasthi;**

**Overview**

The significant expansion in network size to support new paradigms such as cloud computing, IoT (Internet of Things), etc. together with the exponential increase in vulnerabilities has challenged the existing security mechanisms greatly. These challenges have opened many avenues for research in network security. However, while attack graphs play an important role in analyzing vulnerabilities, analyzing large attack graphs itself is a major issue. Therefore, it is necessary to extract only the critical part of the attack graph. Although technologies have been developed for attack path characterization, there is a lack of hybrid technology that can differentiate between similar behavior attack paths. We have proposed a cost-based path characterization technique that takes the attack node's vulnerability complexity into account and significantly reduces the number of vulnerabilities that need to be patched to avoid the major segment of attack graph. Moreover, we have used a real network prototype to validate the performance of the proposed scheme. The proposed scheme works well in cases where some vulnerabilities have similar risk scores. To the best of our knowledge, this is the first time that a cost-effective approach for attack path analysis has been proposed.

# Developing a Smart and Fast Semantic Intrusion Detection System

**PI: Ricardo Calix** *Funding Source: Northrop Grumman*

**Overview**

Intrusion Detection and Prevention Systems (IDS/IPS) serve a pivotal role in securing computer networks. Using machine learning for an Intrusion Detection System is important to stop newattacks that do not have known signatures. The further lowering of the barrier to entry for microprocessor based systems has made it possible to use specialized machine learning coprocessors to improve analysis performance. This grant project proposes a machine learningapproach on a small, low powered embedded system that uses network based features to predict between normal and abnormal network traffic. A hardware based approach using a machine learning.

# CICI: CE: Enhancing Cybersecurity for Broadening Data-Driven Research and Partnerships

**PIs: Sonia Fahmy, Bruno Ribeiro, Xiao Zhu, Busiime Ida Ngambeki, Nicole Key**

**Overview**

As computational and data-driven modeling is leveraged in most disciplines, security of campus cyberinfrastructure is becoming increasingly important not only to protect intellectual property, but also to secure a competitive advantage. Growth in facilities and projects funded by industrial partners and US

government agencies creates new requirements for information security and assurance. Designed for open science, the Purdue campus cyberinfrastructure does not currently have infrastructure in place to support the wide variety of security requirements stemming from the growing volume of aerospace, defense, and industrial projects. This rapid pace of expanding partnerships has outpaced the campus cybersecurity framework, limiting the opportunities for researchers and students. This project brings together engineering and research teams to develop a cyber attack detection and response capability for the Purdue University campus research network.

The project enhances the Purdue campus network to develop a security-assured research network infrastructure, and deploys novel adaptive monitoring tools as part of the campus cyberinfrastructure. Enhanced campus cyberinfrastructure empowers domain scientists to conduct research with heightened security requirements. The project also provides cybersecurity researchers with production network traffic data to develop and evaluate adaptive intrusion detection and response and flow-level anomaly detection. The proposed research will reduce the overhead of network security mechanisms, and improve their effectiveness, agility, and scalability. The project supports cybersecurity education and training by engaging students in deployment and operation of networking hardware and software. It provides students with a unique opportunity to apply their conceptual knowledge to practical situations and better prepares them to be part of the future workforce in cybersecurity. The highly interdisciplinary project team consists of security, networking and data science researchers, cybersecurity educators, campus IT professionals and the engineering community.

## Router Models and Downscaling Tools for Scalable Security Experiments

**PI: Sonia Fahmy**

**Current Students: Ravish Khosla, Wei-Min Yao**

**Overview**

*Funding Source: National Science Foundation (NSF), Northrop Grumman*

A major challenge that researchers face in studying attacks over the Internet is the size of the network to be investigated. For example, a typical Denial of Service (DoS) attack usually takes place over a large portion of the Internet and involves a considerable number of hosts. This can be intractable for testbed experimentation, and even simulation. The goal of this project is twofold: (1) devise solutions to the scalability problem for both network simulation and emulation experiments by partitioning a large network experiment into multiple smaller experiments which are manageable in the simulation or emulation testbed, and (2) develop device-agnostic simulation and emulation models for forwarding devices, such as switches and routers, and design an automated model parameter inference process.

Representative Publications

J. Mirkovic, A. Hussain, S. Fahmy, P. Reiher, R. K. Thomas, “Accurately measuring denial of service in simulation and testbed experiments,” IEEE Transactions on Dependable and Secure Computing (TDSC), volume 6, issue 2, pp. 81-95, Apr/June 2009.

Wei-Min Yao, Sonia Fahmy, "Downscaling Network Scenarios with Denial of Service (DoS) Attacks," In Proceedings of the IEEE Sarnoff Symposium (Invited Session on Security), 6 pages, April 2008.

Roman Chertov, Sonia Fahmy, and Ness B. Shroff, "Fidelity of Network Simulation and Emulation: A Case Study of TCP-Targeted Denial of Service Attacks," ACM Transactions on Modeling and Computer Simulation (TOMACS), volume 19, issue 1, pp. 4:1-4:29, December 2008.

Wei-Min Yao, Sonia Fahmy, "Partitioning Network Testbed Experiments," In Proceedings of IEEE International Conference on Distributed Computing Systems (ICDCS), 11 pp., June 2011.

Project URL: http://www.cs.purdue.edu/homes/fahmy/emist/

## Building Sophisticated Services with Programmable Anonymity Networks

**PIs: Christina Garman, Dave Levin (UMD)**

**Current Students: Arushi Arora**

*Funding Source: National Science Foundation (NSF)*

### Overview

Anonymity systems are critical in achieving free, open communication on today's Internet. In particular, Tor, a popular peer-to-peer anonymous system, has become a staple in resisting online censorship by rogue nations and allowing journalists to safely communicate with sources world-wide.

However, there is a surprisingly narrow set of services that Tor is able to support in a robust fashion. Today, the use of Tor is largely relegated to web proxies and hidden services, and, unfortunately, neither of these applications have the ability to scale to handle dynamic workloads or attacks by automated bots.

Conversely, services on the standard, "non-anonymous" Internet are thriving like never before. Impressive innovations in software-defined networking (SDN), network function virtualization (NFV), content delivery networks (CDNs), and network capabilities have resulted in more robust, scalable, and resilient network services. The present and future Internet is comprised of programmable networks, but there do not exist the basic primitives to achieve such features in anonymous networks.

### Intellectual Merit

Given these trends, the PIs propose an ambitious research agenda towards developing programmable anonymity networks—extensions of Tor that allow users to install and run small snippets of code on Tor routers—and using them to build more sophisticated, more secure anonymous services. The proposed research has four main thrusts:

- *Programmable Tor middleboxes*: Programmable middleboxes are becoming the linchpin in many complex networked systems on the Internet today, as they allow services like load balancers, firewalls, and traffic shapers to be dynamically deployed and migrated. The PIs propose to develop new middleboxlike primitives based on secure enclaves (e.g., SGX) that allow Tor routers to run

small snippets of code on behalf of a third party. If successful, this research will provide a powerful building block that can be used to build a wide range of systems, which we explore in the remaining thrusts.

- *Censorship-resistant hidden services*: One of the fundamental threats to Tor is censorship and deanonymization attacks by routing-capable adversaries. Recent work has made Tor more resilient to censorship, but requires full knowledge of the end-to-end circuit, which is impossible with hidden services. The PIs propose to apply their programmable Tor middleboxes to develop censorship-resistance schemes for hidden services.
- *Hidden-services-based CDN*: Content delivery networks (CDNs) host content on a set of globally replicated servers and protect their customers from large denial of service attacks by stopping attack traffic far from the target's servers. The PIs propose to apply programmable Tor middleboxes to build a CDN from hidden service hosts, demonstrating the capabilities of anonymous middleboxes for dynamic scaling, load balancing, and filtering attack traffic.
- *Decentralized anonymous credentials:* The CDN we propose to build will face the same threats that today's (non-anonymous) CDNs face: that of automated bots. Inherently, anonymity systems cannot tie a connection to a particular user, obviating user or IP-based reputation schemes. The PIs propose to develop anonymous credential schemes that allow users to prove their humanity once and obtain a set of anonymous credentials that can be redeemed at CDNs and our middleboxes for access to resources.

**Broader Impact**

Anonymous communication is a key ingredient to combatting online censorship and suppression of thought and information. Unfortunately, anonymity networks lack the necessary primitives to build sophisticated anonymous systems that are secure against powerful nation-state adversaries. If successful, our proposed research can help improve these trends by: (1) Developing and releasing new tools (and resulting datasets) that help protect existing anonymity networks against powerful attackers and enable new, more sophisticated anonymous systems, and (2) Presenting our findings not only to other researchers, but also to administrators such as NANOG and CISO meetings. Beyond these intellectual impacts, if successful, there will be educational impact, as well; the PIs are dedicated to encouraging women and underrepresented minorities to pursue research, and we anticipate that the goals of this research (free and open communication, protecting journalists, and so on) will entice students to study security who may not have considered it otherwise.

### Representative Publications

Michael Reininger, Arushi Arora, Stephen Herwig, Nicholas Francino, Christina Garman, Dave Levin. "Bento: Bringing Network Function Virtualization to Tor". In ACM CCS Poster Session, 2020.

Michael Reininger, Arushi Arora, Stephen Herwig, Nicholas Francino, Jayson Hurst, Christina Garman, Dave Levin. "Bento: Safely Bringing Network Function Virtualization to Tor". In SIGCOMM 2021.

Provably Avoiding Geographic Regions for Tor's Onion Services. Arushi Arora and Raj Karra (Purdue University); Dave Levin (University of Maryland); Christina Garman (Purdue University). In Financial Cryptography and Data Security 2023.

Project URL: https://barc.cs.purdue.edu/projects/bento.html

# SecureCDN: Providing End-to-End Security in Content Delivery Networks

**PIs: Christina Garman, Dave Levin (UMD)**

### Overview

Content Delivery Networks (CDNs) serve a large and increasing portion of today's web content. Beyond caching, CDNs provide their customers with a variety of services, from load balancing, to content compression and transcoding, to web application firewalls. As web traffic shifts from HTTP to HTTPS, CDNs continue to provide such services by also assuming control of their customers' private keys, thereby breaking a fundamental security principle: private keys must only be known by their owner.

We present the design and implementation of SecureCDN, a reverse caching proxy that uses Intel SGX to preserve the confidentiality of the content provider's private TLS key while stored on the edge server. SecureCDN runs the NGINX webserver in an Intel SGX enclave, while also enabling key CDN services, such as firewalling, local and remote caching, and scriptable configuration. In order to ensure the integrity and, optionally, confidentiality, of any cached content, we also develop a filesystem to extend the enclave's security guarantees to untrusted storage. In its strongest configuration, SecureCDN reduces the knowledge of the edge server to that of a traditional on-path HTTPS adversary. We evaluate the performance of SecureCDN with a series of micro- and macro-benchmarks.

This is ongoing work, in collaboration with the University of Maryland.

# Safeguarding Next-Generation Emergency Services (NG-9-1-1) over Cellular Networks

**PIs: Chunyi Peng, Guan-Hua Tu (Michigan State University)**

### Overview

The next-generation 9-1-1 (NG-9-1-1) services are IP-based systems that supports all types of emergency communications, e.g., voice, video, and text. The globally deployed 5G/4G cellular networks with ubiquitous coverage are the most accessible vehicles for emergency services. However, the security of IP-based cellular emergency services is still largely unexplored. Since 911 services are critical to our society, it is extremely important to eliminate potential vulnerabilities and ensure their security. This project aims to safeguard NG-9-1-1 services over cellular networks from their designs to operations. The project's novelties are to systematically uncover insecure design defects, operational slips, and implementation flaws of cellular emergency services while ensuring sustainable security through a multi-disciplinary approach. The project's broader significance and importance are to extend the state-of-the-art emergency service security research to a new frontier by exploring novel methods that can help perform innovative emergency service modeling, vulnerability analysis, data collection/mining, and testing automation solutions. The project also aims to contribute to nation's future workforce by training students in critical areas such as network security, wireless and mobile systems, and machine learning.

### Representative Publications

Uncovering Insecure Designs of Cellular Emergency Services (911), Yiwen Hu, Min-Yue Chen, Guan-Hua Tu, Chi-Yu

Li, Sihan Wang, Jingwen Shi, Tian Xie, Li Xiao, Chunyi Peng, Zhaowei Tan, Songwu Lu, MobiCom'22, Oct 2022. (Best Community Paper Award Runner-up)

## Formalizing Enterprise Firewall Management with Informal and Elastic Specifications

**PIs: Xiaokang Qiu, Sanjay Rao**

**Current Students: Chenan Wen, Yizhan Qing**

*Funding Source: National Science Foundation (NSF)*

### Overview

Managing enterprise network firewalls is an ad-hoc process today where administrators must extract policy relevant to their enterprises from thousands of natural language vendor documents and tailor them to their unique context. This project aims to achieve formal management of enterprise firewall policy when specifications are informal and incomplete. As the first step, we plan to design a domain specific language to describe elastic firewall specifications and develop a formal semantics for the language so that network operators can interactively describe their specifications and enterprise configurations can be formally and systematically checked.

## Scalable and Resilient Distributed Algorithms for Coordination in Large-Scale Networks

**PI: Shreyas Sundaram** **Current Students: Aritra Mitra, Kananart Kuwaranancharoen**

*Funding Source: National Science Foundation (NSF)*

### Overview

A key challenge in large-scale networked systems is to allow the individual nodes to cooperatively take actions by repeatedly interacting and exchanging information with the other nodes. This is particularly important when the nodes have access to local information, but must take optimal actions that rely on global data. However, large-scale networks also present multiple entry points for attackers to compromise nodes, causing them to behave in unanticipated ways.

In this project, we formulate scalable and lightweight distributed algorithms to allow nodes in large-scale networks to cooperatively take actions, despite malicious behavior by some of the nodes. Our algorithms provide provable safety and performance guarantees for the non-adversarial nodes in the face of worst-case (and possibly coordinated) adversarial behavior by a potentially massive number of attackers. Our algorithms only require each non-adversarial node to interact with its neighbors, and does not require them to know anything about the global network topology. Our research also leads to new metrics for measuring the resilience of networks to attacks, and designing resilient networks. Our algorithms can be applied to the canonical problems of distributed consensus, distributed optimization, and distributed state estimation, among others.

Through this project, we also formulate techniques to design network topologies to enable coordination. We formulate both stochastic and game-theoretic models for the formation of large-scale complex

networks, and identify features of such networks that enable efficient exchange of information and resilience to adversarial behavior.

**Representative Publications**

*H. LeBlanc, H. Zhang, X. Koutsoukos and S. Sundaram, "Resilient Asymptotic Consensus in Robust Networks." IEEE Journal on Selected Areas in Communications: Special Issue on In-Network Computation, vol. 31, no. 4, pp. 766 - 781, Apr 2013.*

*H. Zhang, E. Fata and S. Sundaram, "A Notion of Robustness in Complex Networks." IEEE Transactions on Control of Network Systems, vol. 2, no. 3, pp. 310 - 320, Sept 2015.*

M. Pirani, E. Moradi Shahrivar, B. Fidan and S. Sundaram, "Robustness of Leader-Follower Networked Dynamical Systems" *IEEE Transactions on Control of Network Systems (to appear).*

E. Moradi Shahrivar, M. Pirani and S. Sundaram, "Spectral and Structural Properties of Random Interdependent Networks." *Automatica*, vol. 83, pp. 234 – 242, Sept. 2017.

E. Moradi Shahrivar and S. Sundaram, "The Game-Theoretic Formation of Interconnections Between Networks." *IEEE Journal on Selected Areas in Communications: Special Issue on Game Theory for Networks*, vol. 35, no. 2, pp. 341 - 352, Feb. 2017.

E. Moradi Shahrivar and S. Sundaram, "The Strategic Formation of Multi-Layer Networks." *IEEE Transactions on Network Science and Engineering*, vol. 2, no. 4, pp. 164 - 178, Oct - Dec 2015.

M. Pirani and S. Sundaram, "On the Smallest Eigenvalue of Grounded Laplacian Matrices." *IEEE Transactions on Automatic Control*, vol. 61, no. 2, pp. 509 - 514, Feb. 2016.

S. Sundaram and B. Gharesifard, "Distributed Optimization Under Adversarial Nodes." *IEEE Transactions on Automatic Control*, 2018 (To appear).

D. Senejohnny, S. Sundaram, C. De Persis and P. Tesi, "Resilience against misbehaving nodes in Selftriggered Coordination Networks." Proceedings of the IEEE Conference on Decision and Control, Miami Beach, FL, 2018 (to appear).

K. Kuwaranancharoen and S. Sundaram, "On the Location of the Minimizer of the Sum of Strongly Convex Functions." Proceedings of the IEEE Conference on Decision and Control, Miami Beach, FL, 2018 (to appear).

A. Prasad, H.-L. Choi and S. Sundaram, "Min-Max Tours for Task Allocation to Heterogeneous Agents." Proceedings of the IEEE Conference on Decision and Control, Miami Beach, FL, 2018 (to appear).

A. Mitra and S. Sundaram, "Secure Distributed State Estimation of an LTI system over Timevarying Networks and Analog Erasure Channels." Proceedings of the American Control Conference, Milwaukee, WI, 2018 (to appear).

Project URL: https://engineering.purdue.edu/~sundara2/

# Enabling Detection of Elusive Malware by Going Out of the Box with Semantically Reconstructed View (OBSERV)

**PIs: Dongyan Xu, Xuxian Jiang, George Mason University**

### Overview

There is an alarming trend that elusive malware is armed with techniques that detect, evade, and subvert malware detection facilities of the victim. On the defensive side, a fundamental limitation of traditional host-based anti-malware systems is that they run inside the very hosts they are protecting, making them vulnerable to malware's counter-detection and subversion. To address this limitation, solutions using virtual machine (VM) technologies advocate placing the malware detection facility outside of the protected VM. However, a dilemma exists between these two approaches: The "out of the box" approach gains tamper resistance at the cost of losing the native, semantic view of the host enjoyed by the "in the box" approach. To resolve the above dilemma, a new approach called OBSERV ("Out of the Box with Semantically Reconstructed View") is introduced to achieve the advantages of both camps by reconstructing the semantic internal view of a VM from external, low-level observations. OBSERV enables two exciting malware defense opportunities: (1) malware detection by view comparison and (2) real-time detection and stoppage of kernel-level rootkits. The broader impact of this research is two-fold: (1) It will enhance the trustworthiness and effectiveness of widely deployed anti-malware systems. Moreover, OBSERV is expected to be viewed favorably by the anti-virus software industry because of its support for existing off-the-shelf anti-virus software. (2) Results from this research will lead to the development of education materials for undergraduate and graduate courses and for professional training sessions.

### Representative Publications

Xuxian Jiang, Xinyuan Wang, Dongyan Xu, "Stealthy Malware Detection Through VMM-Based Out-of-the-Box Semantic View Reconstruction", Proceedings of ACM Conference on Computer and Communications Security (CCS 2007), Alexandria, VA, November 2007.

Project URL: http://cairo.cs.purdue.edu/projects

# Virtualization-Enabled Malware Research

**PI: Dongyan Xu**

### Overview

*Funding Source: Microsoft Research and National Science Foundation through the NMI program under grant number OCI-0504261.*

In the battle against Internet malware, we have witnessed increasingly novel features of emerging malware in their infection, propagation, and contamination strategies – examples include polymorphic appearance, multi-vector infection, self-destruction, and intelligent payloads such as self-organized attack networks or mass-mailing. Furthermore, the damages caused by a malware incident can be detrimental and hard to recover (e.g., the installation of kernel-level rootkits). Our research goal is to thoroughly understand key malware behavior such as probing, propagation, exploitation, contamination, and "value-added" payloads. These results will be used to design effective malware detection and defense solutions. To reach this goal,

we realize that effective malware experimentation tools and environments are lacking in current malware research. By leveraging and extending virtualization technology, we propose to develop a virtualization-based integrated platform for the capture, observation, and analysis of malware. The platform consists of two parts: The front-end of the platform is a virtual honey farm system called Collapsar, which captures and contains malware instances from the real Internet. The back-end of the platform is a virtual playground environment called vGround, where the captured malware instances are unleashed to run while remaining completely isolated from the real Internet. Using this integrated platform, security researchers will be able to observe and analyze various aspects of malware behavior as well as to evaluate corresponding malware defense solutions, with high fidelity and efficiency.

### Representative Publications

Ryan Riley, Xuxian Jiang, Dongyan Xu, "An Architectural Approach to Preventing Code Injection Attacks", Proceedings of IEEE/IFIP International Conference on Dependable Systems and Networks (DSN-DCCS 2007), Edinburgh, UK, June 2007.

Xuxian Jiang, Dongyan Xu, "Collapsar: A VM-Based Architecture for Network Attack Detention Center," Proceedings of the 13th USENIXSecurity Symposium, San Diego, CA, August 2004.

Project URL: http://cairo.cs.purdue.edu/projects

## Big Data Security Analyses

### PI: Baijian Yang

### Overview

Study how to better use macine learning to quickly identify events in big data environment.

### Representative Publications

B. Yang and T. Zhang, "A Scalable Feature Selection and Model Updating Approach for Big Data Machine Learning", to appear IEEE SmartCloud 2016, NYC, USA, 2016
T. Zhang and B. Yang, "Big Data Dimension Reduction using PCA", to appear IEEE SmartCloud 2016, NYC, USA, 2016

S. Nanda, F. Zafari, C. DeCusatisy, E. Wedaaz and B. Yang , "Predicting Network Attack Patterns in SDN using Machine Learning Approach", IEEE NFV-SDN 2016, Palo Alto, CA, USA

Zhang, T., Yang, B. (2016). Box-Cox Transformation in Big Data. Technometrics. www.tandfonline.com/loi/utch20

Zhang, T., & Yang (2018), B. Dimension reduction for big data. Statistics and Its Interface, 11(2), 295-306.

Ryu, S.-H.G, & Yang, B. Comparison of Machine Learning Algorithms and Their Ensembles for Botnet Detection. Dekalb, IL: International Conference of Information and Computer Technology, 2018.

Zhang, T., & Yang, B. (2017). An exact approach to ridge regression for big data. Computational Statistics, 32(3), 909-928.

# Development of a Secure and Privacy-Preserving Workflow Architecture for Dynamic Data Sharing in Scientific Infrastructures

**PI: Xukai Zou, Huanmei Wu and Saptarshi Purkayastha**

*Funding Source: National Science Foundation (NSF)*

## Overview

Scientific cyberinfrastructures embrace collaborative workflows where users can access and share heterogeneous data and computing resources to perform research and education tasks, which catalyze scientific discovery. One such cyberinfrastructure, JetStream, is the first production cloud funded by the NSF for general-purpose science and engineering research and education. Although Jetstream provides basic data storage security and web authentication, its security features do not satisfy the strict requirements involving sensitive data, such as healthcare data with protected health information (PHI). This project will build a secure, holistic and resilient cybersecurity architecture on JetStream so that collaborative research and education projects can share PHI securely between its users.

The secured infrastructure will provide comprehensive multi-level protection for the PHI and its workflows through user authentication, fine-tuned data access control, confidentiality, integrity, and traceability. The project will implement advanced security techniques, such as role-wise passwordless authentication and authorization, cryptography-based hierarchical access control, dual-level key management, and secure digital provenance or blockchain-based integrity protection. By employing these, JetStream VMs will be able to guarantee the security, privacy, and integrity of scientific workflows and associated data, thus protecting data and computing resources from internal and external attacks. When applied to healthcare and life-science cyberinfrastructures, it will enable sensitive health data to be shared securely, which is an essential requirement for accelerating life science research. The project will promote the use of real clinical data in training to produce enormous educational impacts. The developed secure architecture is generic and applicable to other data and resource sharing environments.

# Other Security Research

## Improving the Security and Usability of the Wear OS Permission Model

**PIs: Berkay Celik, Antonio Bianchi**

*Funding Source: Google, Inc. (partial funding)*

**Overview**

Google's Wear OS is a version of Android's operating system specifically designed to manage wearable devices, such as smartwatches and other wearables. Normally, Wear OS apps have the ability to access potentially sensitive information, such as the device's location and is controlled by a permission system. Specifically, users are asked at run-time whether they want to allow a Wear OS app to access a specific piece of sensitive information. Through permissions, Android allows the user to select whether to allow or deny sensitive information access to the app.

However, with potentially confusing permissions windows popping up in a dialog box, it's possible for the user to inadvertently choose options that send location data to the Wear OS app. This reveals three fundamental issues; poor usability, poor user understanding, and unclear security.

We plan to perform what would be the first systematic analysis of the interaction of the Android/Wear OS permission models.

## Investigating and Understanding Digital Bill of Materials

**PI: Christina Garman**

**Current Students: Arushi Arora**

*Funding Source: Idaho National Laboratory*

**Overview**

Managing the integrity, authenticity, and reliability of critical systems including Managing the integrity, authenticity, and reliability of critical systems including high priority operational technology (OT) components is becoming increasingly im-portant across various government sectors, as evidenced by the recent executive order on cybersecurity, EO 14028. Digital supply chains that sustain these critical in-frastructures are growing frequently diverse and complex, resulting in variations in the overall cybersecurity risk for energy systems. The transparency between vendors andasset owners concerning the hardware and software components of their adopted toolsalong with the security vulnerabilities within them have the potential to aid the sector'sability to mitigate risk.

The Digital Bill of Materials (DBoM) is a proposed solution to this obstacle whichallows sharing Software

Bill of Materials (SBoM), Hardware Bill of Materials (HBoM), and attestations across vendors through a set of supported repositories. Current imple-mentations of DBoM introduce problems encompassing sharing of such crucial and sensitive (some of it may also be confidential) information in a secure and privacy-preserving approach. One such obstacle that a vendor may encounter is while sharing security-relevant information regarding a proprietary subcomponent producedby a third party and contained in the vendor's software. Moreover, sharing softwarevulnerabilities in the digital supply chain may allow a malicious entity access to such sensitive information even before the vulnerability is patched. We wish to seek an-swers to questions like how could an entity check for the presence of software compo-nents (and subcomponents) or vulnerabilities anonymously and securely or how could a party trust an SBoM or HBoM entry. During the course of this project, we aim to de-fine DBoM threat models, identify security and privacy concerns that may arise while sharing such critical information, and attempt to propose solutions for the same.

### Representative Publications

Arora, Arushi, Virginia Wright, and Christina Garman. "Strengthening the Security of Operational Technology: Understanding Contemporary Bill of Materials." Journal of Critical Infrastructure Policy• Volume 3.1 (2022).

## Removing the Human Element: Securing Deployed Cryptographic Systems through the use of Cryptographic Automation

**PI: Christina Garman**

**Current Students: Yongming Fan, Jacob White**

*Funding Source: National Science Foundation (NSF)*

### Overview

Cryptography has shown itself to be invaluable in everyday life, especially as more and more devices and interactions are moving to the online world. Whether it is browsing the web, making a purchase, or sending a message to a friend, cryptography is everywhere. Despite the fact that users (often unknowingly) rely on the security of systems that use cryptography, recent years have seen a number of serious vulnerabilities in the cryptographic pieces of systems, some with large consequences. These have been caused by various problems, including poor designs, difficulty of implementation, and use (or misuse) of (in)secure primitives. There is a common denominator in all of these problems: the human element. Many of the errors that are found when analyzing these insecure systems could have been prevented if both designers and software engineers had better tools to help them navigate the complex cryptographic space. Cryptographic automation is a relatively new and promising area that is designed to help solve many of these issues and make developing secure systems far easier and less error-prone, even for a non-expert. This project focuses on removing the human element from the deployment and analysis of cryptographic systems. Through the use of cryptographic automation and the development of tools, the project's aim is to make it easier to design and securely deploy new and complex cryptographic systems while preventing insecurities from occurring in such systems. Additionally, the project contains an education plan designed to help make cryptography more accessible to a broader audience. The creation of the Midwest Women in

Computer Security Workshop, as well as the project's goal to not just develop but also disseminate tools, will allow more students of all ages, and more software engineers, to explore cryptography and computer security, instead of being intimidated or afraid of it.
The project has three main thrusts. The core of the project centers around the first thrust of building tools to aid in the deployment of complex cryptography. This will principally focus on automating the end-to-end development of zero-knowledge proof code, from expressing the proof statement to realizing the implementation, with additional applications to anonymous credentials. The second thrust focuses on automating the discovery of cryptographic vulnerabilities in applications that use zkSNARKs, a popular zero-knowledge proof instantiation. This thrust will leverage fuzzing to help both programmers and end users detect inconsistencies and errors in existing, already deployed zkSNARK circuits and applications. The third thrust works to automate the discovery and identification of modern cryptographic algorithms and techniques in both traditional as well as heavily obfuscated binaries, through a novel combination of various dynamic analysis and machine-learning based approaches. If successful, the combination of these three thrusts will, for expert and non-expert developers alike, make it both easier to discover the use of cryptography and potentially vulnerable algorithms in existing systems as well as design and securely deploy new and complex cryptographic systems while preventing these insecurities from happening.

## SGX.Fail: How Secrets Get eXtracted

**PI: Christina Garman** **Current Students: Alex Seto**

### Overview

Intel's Software Guard Extensions (SGX) promises an isolated execution environment, protected from all software running on the machine. As such, numerous works have sought to leverage SGX to provide confidentiality and integrity guarantees for code running in adversarial environments. In the past few years however, SGX has come under heavy fire, threatened by numerous hardware attacks. With Intel repeatedly patching SGX to regain security while consistently launching new (micro)architectures, it is increasingly difficult to track the applicability of various attack techniques across the SGX design landscape.

Thus, in this project we set out to survey and categorize various SGX attacks, their applicability to different SGX architectures, as well as the information leaked by them. We then set out to explore the effectiveness of SGX's update mechanisms in preventing attacks on real-world deployments. Here, we study two commercial SGX applications. First, we investigate the SECRET network, an SGX-backed blockchain aiming to provide privacy-preserving smart contracts. Next, we also consider PowerDVD, a UHD Blu-Ray Digital Rights Management (DRM) software licensed to play discs on PCs. We show that in both cases vendors are unable to meet security goals originally envisioned for their products, presumably due to SGX's long update timelines and the complexities of a manual update process. This in turn forces vendors to make difficult security/usability trade offs, resulting in security compromises.

### Representative Publications

Stephan van Schaik, Alex Seto, Thomas Yurek, Adam Batori, Bader AlBassam, Christina Garman, Daniel Genkin, Andrew Miller, Eyal Ronen, Yuval Yarom. "SoK: SGX.Fail: How Stuff Get eXposed".

Project URL: https://sgx.fail/

# SNARKProbe: An Automated Security Analysis Framework for zkSNARK Implementations

**PI: Christina Garman  Current Students: Yongming Fan**

## Overview

With the recent growing interest in privacy-enhancing technologies, we are seeing a complementary growth in the desire to build and deploy complex cryptographic systems that involve techniques like zero-knowledge proofs. Of these, general purpose proof systems like zkSNARKs have seen the most interest, due to their small proof size, fast verification, and expressiveness. Unfortunately, as we have seen with many areas of cryptography, guaranteeing correct implementations can be tricky, as the protocols themselves are complicated and often require substantial low-level manual effort to achieve maximum performance. To help with this problem, and gain better assurances about the correctness and security of already implemented zkSNARK protocols and applications, we design and build SNARKProbe, an automated security analysis framework for zkSNARKs that can scan R1CS-based libraries and applications to detect various issues, such as edge case crashing, cryptographic operation errors, and/or inconsistencies with protocol descriptions. SNARKProbe leverages a variety of analysis techniques, including fuzzing and SMT solvers. We test the performance of SNARKProbe on a variety of different experimental parameters to demonstrate its practicality and reasonable runtime, and we also evaluate its ability to find potential inconsistencies and errors in implementations.

# Policy, Law, and Management

## Big Data Ethics: detecting bias in data collection, algorithmic discrimination and "informed refusal"

**PIs: Chris Clifton, Kendall Roark, Daniel Kelly**

*Funding Source: Mellon Foundation*

### Overview

We are increasingly seeing evidence of discriminatory outcomes from data-based decision making. Yet whether due to lack of a nuanced understanding about how big data is collected and how algorithms work, or due to the lack of transparency on the part of data producers and aggregators (or both) the ability for civil society to meaningfully engage in the governance of data collection and use is severely limited. Much debate has occurred in the realm of privacy protection and the ways in which machine learning can reproduce existing bias, from minor but insidious events such as the disparity in web advertising based on gender or race (Sweeney 2013; Datta 2015), imposing additional barriers to access on groups that do not "match the norm" such as Facebook requiring proof of identity for Native Americans (Sanburn 2015), to obvious life-changing outcomes such as predictive models for sentencing guidelines (Angwin 2016). However, there has been less emphasis on the ways in which emerging data collection methods and machine learning on large volumes of data themselves can introduce novel forms of impact on individuals' lives, heightening or introducing new avenues for institutional discrimination. There has also been little attention paid to the difference between legal constraints around individual privacy and the potential harm from aggregate impacts on broader society in the era of data-driven decision making. What is clear is that such data-based decisions at scale have the potential to discriminate in unintended and difficult to detect ways, further blocking efforts to achieve an equitable and just society.

We are addressing these grand challenges through a multidisciplinary study of the ethical issues involved in the use of big data and predictive algorithms to make decisions affecting individuals. We will assemble a concrete set of cases, and use these to define the more general problem or problems that arise. Some of these cases will come from existing studies in data-driven discrimination (Sweeney 2013; Datta 2015; Angwin 2016). Others will involve historic discrimination data. We will also look at public data such as the NIJ Crime Forecasting Challenge (NIJ 2016), and public social media, where differences between groups may include distinctions based on personal preference as well as distinctions based on group stereotyping.

# Purdue University's Computer and Information Technology program creates framework to support implementing data governance in small and medium enterprises.

**PI: Chad Laux, John Springer** **Current Students: Manal Alduraibi**

## Overview

Purdue's Computer and Information Technology PhD program investigated the effects of the Lean Six Sigma (LSS) quality approach in supporting the implementation of data governance and data privacy in small and medium enterprises (SMEs). Working in conjunction with Indiana's Indiana Executive Council on Cybersecurity (IECC), and under the guise of Dr. Chad Laux, the research team created a framework for implementing and improving operations in data governance, utilizing Lean and Six Sigma tools.

Based on the findings of this study, it is recommended that SMEs implement Lean Six Sigma principles to improve their data governance systems and overall operational performance. To do so effectively, SME leaders should use the framework developed in this study as a guide, taking into account the challenges and opportunities that arise during implementation. Additionally, SMEs should consider the potential benefits of using LSS tools such as process mapping, risk analysis, mistake proofing, and benchmarking to support their data governance efforts. It is also recommended that SMEs invest in training and education programs to ensure that their employees have the necessary skills and knowledge to implement and maintain effective data governance systems.

The impact of this study is a practical approach to serve the approximately 534,000 business and 1.2 million employees that comprise Indiana's small and medium enterprise business profile.

## Representative Publications

Manal Alduraibi (2023). Investigating the Impact of Lean Six Sigma Principles on Establishing and Maintaining Data Governance Systems in Smes: An Exploratory Study Using Grounded Theory and Ism Approach.

Project URL: http://Program link: https://polytechnic.purdue.edu/departments/computer-and-information-technology

## Assessing Security for Organizations Dealing with At-Risk Populations

**PIs: Gene Spafford, Kelley Misata (former PhD)**

### Overview

Organizations that deal with at-risk populations -- including battered spouses, human trafficking victims, etc -- are themselves at risk. They often do not have resources, funding, or personnel to adequately assess their cyber risks or to protect against them. Their clients and staff, meanwhile, are at significant risk, including possibly for kidnapping and assault.

This project is intended to provide a baseline understanding of the needs and capabilities of these organizations, and develop a strategy to help them improve their cyber protection posture.

The result is a non-profit that provides consulting and education for other non-profits.

### Representative Publications

Gap Analysis Identifying the Current State of Information Security within Organizations Working with Victims of Violence by Kelley Misata, PhD Dissertation.

Project URL: http://See https://sightlinesecurity.org/

## CICI: RDP: Supporting Controlled Unclassified Information with a Campus Awareness and Risk Management Framework

**PIs: Baijian Yang, Preston M Smith**

### Overview

Protecting Controlled Unclassified Information (CUI) is mandated by the executive order 13356, and today is required for research in sectors such as defense and aerospace. Regulatory requirements for research will increase, with CUI regulations covering categories including Agriculture, Financial, Legal Records, and Business information. When combined with existing regulations already seen by universities, such as HIPAA, and the European Union's GDPR, a well-defined and consistent framework for working with regulated data is critical for institutions of higher education. This project describes a cost-effective ecosystem (REED+) to manage regulated data that meets the compliance requirements found in a campus environment. The REED+ framework integrates NIST SP 800-171 and other related NIST publications as the foundation of the framework. This framework serves as a standard for campus IT to align with security regulations and best practices, and create a single process for intake, contracting, and facilitate easy mapping of controlled research to CI resources for the sponsored programs office, human subjects office, and export control office. The framework allows researchers to experience faster intake of new funded projects and be more competitive for research dollars. Using student-developed training materials and instruction, researchers, administrators, and campus IT are now able to more clearly understand previously complicated data security regulations affecting research projects. The ecosystem developed from this project enables new partnerships with government agencies, and industry partners from the defense, aerospace, and life science sectors. Experiences and best practices in providing cyberinfrastructure and security awareness developed from this collaboration are documented and shared with the broader CI and campus community through conferences, journals and workshop.

# Prevention, Detection and Response

## Subtle Adversarial Intrusion Detection with SONAR Software

**PI: Hany Abdel-Khalik**

**Current Students: Tyler Lewis, Yeni Li**

*Funding Source: Idaho National Laboratory*

### Overview

The Signal-Oriented Network Anomaly Recognition (SONAR) tool is an extension of Idaho National Laboratory's (INL's) Risk Analysis Virtual Environment (RAVEN) framework focused on data-level intrusion detection. SONAR aims to expand the already-widespread capabilities of the RAVEN framework by providing additional tools for both data analysis and intrusion detection that are fully compatible with downstream analyses performed by RAVEN.

SONAR utilizes a decomposition-based approach to categorize unknown signals based on their similarity to genuine articles in the so-called feature space. The tool has been designed with flexibility in mind, so the type of distance metric utilized, e.g., cosine distance, Euclidean distance, etc., and manner of decomposition, e.g., dynamic mode decomposition, Fourier decomposition, etc., are flexible to the needs of the user. The advantage of intrusion detection in the feature space is that subtle variations produce outsized differences in high-ranking characteristics that are not otherwise detectable. By considering low-distance signals as 'genuine', SONAR has shown remarkable consistency in identifying subtle intrusions in a variety of experimental and simulated datasets; the various intrusions have been modelled as a modified form of the well-known Triangle Attack to generate a well-hidden data perturbation within the standard behavior of the data that would typically go undetected in real-world and industrial systems. SONAR, as a tool, also provides numerous data analysis capabilities, including visualization, decomposition, and a parameter sweeping protocol, in order to optimize the effectiveness of intrusion detection; this capability allows for further specification allowing unknown data samples to be characterized with a high degree of certainty.

### Representative Publications

Yeni Li, Arvind Sundaram, Hany S. Abdel-Khalik & Paul W. Talbot, "Real-Time Monitoring for Detection of Adversarial Subtle Process Variations", Nuclear Science and Engineering, October 2021.

Yeni Li, Paul W. Talbot & Hany S. Abdel-Khalik, "A Novelty Detection Workflow for Nuclear System Monitoring", 2022 ANS Winter Meeting, Pheonix, AZ.

# Behavioral and Game-Theoretic Security Investments in Interdependent Systems

**PIs: Saurabh Bagchi, Shreyas Sundaram, Timothy Cason**

*Funding Source: National Science Foundation (NSF)*

## Overview

Modern cyber-physical systems (CPS) are increasingly facing attacks by sophisticated adversaries. These attackers are able to identify the susceptibility of different targets in the system and strategically allocate their efforts to compromise the security of the network. In response to such intelligent adversaries, the operators (or defenders) of these systems also need to allocate their often limited security budget across many assets to best mitigate their vulnerabilities. This has led to significant research in understanding how to better secure these systems, with game-theoretical models receiving increasing attention due to their ability to systematically capture the interactions of strategic attackers and defenders.

In the context of large-scale interdependent systems, adversaries often use stepping-stone attacks to exploit vulnerabilities within the network in order to compromise a particular target. Such threats can be captured via the notion of attack graphs that represent all possible paths that attackers may have to reach their targets within the CPS. The defenders in such systems are each responsible for defending some subset of the assets with their limited resources. In much of the existing literature, the defenders and attackers are modeled as fully rational decision-makers who choose their actions to maximize their expected utilities. However, a large body of work in behavioral economics has shown that humans consistently deviate from such classical models of decision-making seminal model capturing such deviations is prospect theory (introduced by Kahneman and Tversky in 1979), which shows that humans perceive gains, losses, and probabilities in a skewed (nonlinear) manner, typically overweighting low probabilities and underweighting high probabilities.

We model the behavioral biases of human decision-making in securing interdependent systems and show that such behavioral decision-making leads to a suboptimal pattern of resource allocation compared to non-behavioral (rational) decision-making. We provide empirical evidence for the existence of such behavioral bias model through a controlled subject study with 145 participants. We then propose three learning techniques for enhancing decision-making in multi-round setups. We illustrate the benefits of our decision-making model through multiple interdependent real-world systems and quantify the level of gain compared to the case in which the defenders are behavioral. We also show the benefit of our learning techniques against different attack models.

Representative Publications

Mustafa Abdallah, Parinaz Naghizadeh, Ashish R. Hota, Timothy Cason, Saurabh Bagchi, and Shreyas Sundaram, “Behavioral and Game-Theoretic Security Investments in Interdependent Systems Modeled by Attack Graphs,” in IEEE Transactions on Control of Network Systems (TCNS), accepted for publication in a future issue, pp. 1-12, April 2020.

Mustafa Abdallah, Parinaz Naghizadeh, Timothy Cason, Saurabh Bagchi, and Shreyas Sundaram. "Protecting assets with heterogeneous valuations under behavioral probability weighting." In 2019 IEEE

58th Conference on Decision and Control (CDC), pp. 5374-5379, December 11-13, 2019.

Mustafa Abdallah, Parinaz Naghizadeh, Ashish Hota, Timothy Cason, Saurabh Bagchi, and Shreyas Sundaram, "The Impacts of Behavioral Probability Weighting on Security Investments in Interdependent Systems," At the American Control Conference (ACC), pp 5260-5265, July 10-12, 2019, Philadelphia, PA.

Aritra Mitra, John A. Richards, Saurabh Bagchi, and Shreyas Sundaram, "Resilient Distributed State Estimation with Mobile Agents: Overcoming Byzantine Adversaries, Communication Losses, and Intermittent Measurements," in Springer "Autonomous Robots", vol. 43, no. 3, pp. 743-768, March 2019.

Daniel Woods, Mustafa Abdallah, Saurabh Bagchi, Shreyas Sundaram, Timothy Cason, "Network Defense and Behavioral Biases: An Experimental Study," Journal of Experimental Economics, August 2020.

Project URL: https://engineering.purdue.edu/dcsl/project-new/project-1/

## An MTD-based Self-Adaptive Resilience Approach for Cloud Systems

**PIs: Bharat Bhargava, Norm Ahmed at AFRL, Jason Kobes at NGC**

### Overview

Advances in cloud computing have made it a feasible and cost-effective solution to improve the resiliency of enterprise systems. However, the replication approach taken by cloud computing to provide resiliency leads to an increase in the number of ways an attacker can exploit or penetrate the systems. This calls for designing cloud systems that can accurately detect anomalies and dynamically adapt themselves to keep performing mission-critical functions even under attacks and failures. In this paper, we propose a self-adaptive resiliency approach for cloud enterprise systems that employs a live monitoring and moving target defense based approach to automatically detect deviations from normal behavior and reconfigure critical cloud processes through software-defined networking to mitigate attacks and reduce system downtime. The proposed solution is promising to present a unified framework for resilient cloud systems. Keywords-moving target defense; resiliency; adaptability; cloud security.

### Representative Publications

IEEE Transactions on Cloud Computing, IEEE Conference on Cloud Computing, ACM CCS conference and workshop

Project URL: https://www.cs.purdue.edu/homes/bb/#research

# Autonomous Aggregate Data Analytics in Untrusted Cloud

**PI: Bharat Bhargava, P Goyal in IIT India, Jason Kobes in NGC**

### Overview

Intelligent Autonomous Systems (IAS) are highly reflexive and very cognizant about their limitations and capabilities, interactions with neighboring entities, as well as the interactions with its operational environment. IAS should be able to conduct data analytics and update policies based on those analytics. These tasks should be performed autonomously i.e. with limited or no human intervention. In this paper, we introduce advanced aggregate analytics over untrusted cloud and autonomous policy updates as a result of those analytics. We will be using Active Bundle (AB), a distributed self- protecting entity, wrapped with policy enforcement engine as our implementation service. We propose an algorithm that can enable individual ABs to grant or limit permissions to their AB peers and provide them with access to anonymized data to conduct analytics autonomously. When these processes take place, ABs do not need to rely on policy enforcement engine every time, which increases scalability. This workflow also creates an AB environment that is decentralized, privacy- preserving, and autonomous.

### Representative Publications

Submitted to IEEE Cloud

Project URL: https://www.cs.purdue.edu/homes/bb/#research

# ConFoc: Content-Focus Protection Against Trojan Attacks on Neural Networks

**PI: Bharat Bhargava**

### Overview

Deep Neural Networks (DNNs) have been applied successfully in computer vision. However, their wide adoption in image-related applications is threatened by their vulnerability to trojan attacks. These attacks insert some misbehavior at training using samples with a mark or trigger, which is exploited at inference or testing time. In this work, we analyze the composition of the features learned by DNNs at training. We identify that they, including those related to the inserted triggers, contain both content (semantic information) and style (texture information), which are recognized as a whole by DNNs at testing time. We then propose a novel defensive technique against trojan attacks in the context if image classification, in which DNNs are taught to disregard the styles of inputs and focus on their content only to mitigate the effect of triggers during the classification. The generic applicability of the approach is demonstrated in the context of a traffic sign and a face recognition application. Each of them is exposed to a different attack with a variety of triggers. Results show that the method reduces the attack success rate significantly to values < 1> as well as improving the initial accuracy of the models with both benign and adversarial data.

# Hunting for Insider Threats Using LSTM-based Anomaly Detection

PI: Bharat Bhargava

### Overview

Insider threats are one of the most difficult problems to solve, given the privileges and information available to insiders to launch different types of attacks. Current security systems can record and analyze sequences from a deluge of log data, potentially becoming a tool to detect insider threats. The issue is that insiders mix the sequence of attack steps with valid actions, reducing the capacity of security systems to programmatically detect the attacks. To address this shortcoming, we introduce LADOHD, an anomaly detection framework based on Long-Short Term Memory (LSTM) models, which learns the expected event patterns in a computer system to identify attack sequences even when attacks span for a long time. The applicability of the framework is demonstrated on a dataset of 38.9 million events collected from a commercial network of 30 computers over twenty days and where a 4-day long insider threat attack occurs. Results show that LADOHD outperforms the anomaly detection system used to protect the commercial network with a True Positive Rate of 97.29% and False Positive Rate of 0.38%. Experiments also show that LSTMs have higher prediction precision in variable-length sequences than methods like Hidden Markov Models, a crucial requirement in sequence-analysis-based anomaly detection techniques.

# Incremental Learning Through Graceful Degradations in Autonomous Systems

PIs: Bharat Bhargava, Jason Kobes at NGC

### Overview

Intelligent Autonomous Systems (IAS) are highly cognitive, reflexive, multitasking, trustworthy (secure as well as ethical), and rich in knowledge discovery. IAS are deployed in dynamic environments and connected with numerous devices of different types, and receive large sets of diverse data. They often receive new types of raw data that was not present in either training or testing data sets thus they are unknown to the learning models. In a dynamic environment, these unknown data objects cannot be ignored as anomalies. Hence the learning models should provide incremental guarantees to IAS for learning and adapting in the presence of unknown data. The model should support progressive enhancements when the environment behaves as expected or graceful degradations when it does not. In the case of graceful degradations, there are two alternatives: (1) weaken the acceptance test of data object (operating at a lower capacity) or (2) replace primary system with a replica or an alternate system that can pass the acceptance test. In this paper, we provide a combinatorial design— MACROF configuration—built with balanced incomplete block design to support graceful degradations in IAS and aid them to adapt in dynamic environments. The architecture provides stable and robust degradations in unpredictable operating environments with limited number of replicas. Since the replicas receive frequent updates from primary systems, they can take over primary system's functionality immediately after an adverse event. We also propose a Bayesian learning model to dynamically change the frequency of updates. Our experimental results show that MACROF configuration provides an efficient replication scheme to support graceful degradations in autonomous systems.

Project URL: https://www.cs.purdue.edu/homes/bb/#research

# Machine Learning Models to Enhance the Science of Cognitive Autonomy

**PIs: Bharat Bhargava, Jason Kobes**

### Overview

IAS include software systems that are capable of automatic reconfiguration, autonomous vehicles, network of sensors with reconfigurable sensory platforms, and an unmanned aerial vehicle (UAV) respecting privacy by deciding to turn off its camera when pointing inside a private residence. Research is needed to build systems that can monitor their environment and interactions, learn their capability as well limitations, and adapt to meet the mission objectives with limited or no human intervention. The systems should be fail-safe and should allow for graceful degradations while continuing to meet the mission objectives. In this paper, we propose new methodologies and workflows, and survey the existing approaches and new ones that can advance the science of autonomy in smart systems through enhancements in real- time control, auto-reconfigurability, monitoring, adaptability, and trust.
Project URL: https://www.cs.purdue.edu/homes/bb/#research

# Quantitative Evaluation of Extensive Vulnerability Set Using Cost Benefit Analysis

**PIs: Bharat Bhargava, U. Bansal, G. Sikka, L. K. Awasthi**

### Overview

The significant expansion in network size to support new paradigms such as cloud computing, IoT (Internet of Things), etc. together with the exponential increase in vulnerabilities has challenged the existing security mechanisms greatly. These challenges have opened many avenues for research in network security. However, while attack graphs play an important role in analyzing vulnerabilities, analyzing large attack graphs itself is a major issue. Therefore, it is necessary to extract only the critical part of the attack graph. Although technologies have been developed for attack path characterization, there is a lack of hybrid technology that can differentiate between similar behavior attack paths. We have proposed a cost-based path characterization technique that takes the attack node's vulnerability complexity into account and significantly reduces the number of vulnerabilities that need to be patched to avoid the major segment of attack graph. Moreover, we have used a real network prototype to validate the performance of the proposed scheme. The proposed scheme works well in cases where some vulnerabilities have similar risk scores. To the best of our knowledge, this is the first time that a cost-effective approach for attack path analysis has been proposed.

### Representative Publications

U. Bansal, G. Sikka, L. K. Awasthi and B. Bhargava, "Quantitative Evaluation of Extensive Vulnerability Set Using Cost Benefit Analysis," in *IEEE Transactions on Dependable and Secure Computing*, vol. 21, no. 1, pp. 298-308, Jan.-Feb. 2024, doi: 10.1109/TDSC.2023.3253121.

Project URL: https://ieeexplore.ieee.org/document/10061343

# Scalable Learning Through Error-correcting Codes based Clustering in Autonomous Systems

**PI: Bharat Bhargava**

### Overview

Intelligent Autonomous systems (IAS) continuously receive large streams of diverse data from numerous entities operating and interacting in their environment. It is vital that the learning models in IAS to scale up to the new and unknown data items that were not present in the training or testing datasets. Scalable learning is nothing but a method to achieve maximum classification without rejecting any unknown data item as anomalies. In this paper, we present Perfect Error- correcting Codes (PEC) clustering technique to approximate the classes of multi-feature data items by reversing standard forward error correction coding. Approximating classes problems generally arise in information systems that are processing fuzzily cataloged data items. These data can be classified by applying binary vectors to their corresponding features (1: feature is present or 1: feature is absent) to obtain message words. These code words can be used as cluster centers. In PEC clustering, binary vectors of 23 bits are mapped into code words (labels or indices) of 12 bits. Two binary vectors with the Hamming distance of 2 will have a few common labels thus classified accordingly. PEC clustering has 2 23 code word space, which makes it ideal for scalability in clustering of thousands of categories. With reasonable redundancy, the clustering can be accomplished in O ( N ) time. In addition, we present an information processing model for on-the-fly processing of data streams with multi-processor pipeline: Read, Analyze, and Toggle (RAT) model.

Project URL: https://www.cs.purdue.edu/homes/bb/#research

# Developing New Mechanisms to Enable Open Source Software (OSS) Supply Chain Transparency

**PIs: Sabine Brunswicker, Santiago Torres**

**Current Students: Sahithi Kasim**

### Overview

Open source software (OSS) has significantly increased the complexity of the software supply chain in terms of source code development, building and packaging, re-configuration and re-packaging (for instance, containerizing), and deployment. Nowadays, practically every software developer reuses and integrates packages from openly accessible OSS products when moving through this chain. Significant benefits include reduced development times and lower costs when reusing OSS packages. However, OSS supply chains risk malicious code deployment, attacks on source code repositories, and unsuccessful packaging procedures. In fact, according to the most recent Symantec Internet Security Threat Reports, the complexity of the OSS-dependent software supply chains has led to a cumulative increase in compromised software products of 750%. The technological and organizational dependency across various products and factors grows as more and more open-source software packages are reused and repackaged throughout supply chains. OSS supply chain interdependencies and the risk they pose are highly opaque, even though transparency is the core value of OSS development. Furthermore, the existing infrastructure

does not disclose the risks' sources and consequences for other players in the OSS supply chain ecosystem.

In order to mitigate attacks of this nature, this project focuses on developing data science mechanisms to detect and prevent software supply chain attacks. This requires us to identify various data sources as well as their relationships to better understand threat indicators, the supply chain's vulnerable surface and the behavioral patterns trying to exploit these.

Project URL: https://github.com/TSELab/guac-alytics

## Developing Software Sensors for Digital Twin based Cybersecurity

PI: Berkay Celik

*Funding Source: Rolls-Royce*

### Overview

Digital twins are a virtual copy of physical devices and are used to create digital models that update and change alongside their real-life counterparts. In this effort, we propose to develop digital twins of sensors called software sensors as the virtual backup of the corresponding physical devices. The emergence of computationally powerful physical systems, particularly deployed in mission-critical systems, allows us to install complete monitoring and recovery software components. Taking advantage of this fact, a software sensor will continuously predict the readings of the corresponding physical sensor.

## System Events and Network Traffic Generation for Realistic Cyber Experimentation

### PI: Berkay Celik

### Overview

Virtual testbeds are critical to understanding threats to computer systems and evaluating potential defenses. They are used to construct experimentation in a controlled environment that requires recreating attack scenarios to reason about the nature of attacks more precisely. While such scenarios hold significance for the community, lack of semantically rich reconstructions of real-world attack scenarios undermines realism, and could potentially lead to overly optimistic conclusions, and defenses ineffective in practice. In this work, we focus on developing tools to generate application and network layer semantics that provides a basis for prudent modeling of benign and malicious actors. We design and implement attacks on single and multiple hosts that exploit different vulnerabilities through their APT campaigns reports into the SOL4CE platform. Such attacks enable us to emulate realistic attack behaviors present activities in system events and network traffic. The developed tools are used to emulate users that encompass computational models of human behavior during attack execution. Accomplishing this task demonstrates the efficacy and breadth of our methods in identifying the artifacts after blending legitimate traces with attack traces, such as system logs and explicit information flows including network communication. Our proposed project helps improve the ability of the SOL4CE platform in threat modeling. Through this effort, we provide SOL4CE platform users a means of executing attack scenarios

feasible for realistic deployments and identifying potential threats with system and network data on the end hosts. Overall, the outcome of this effort fosters attack scenario designs and stymie realistic cyber experimentation.

## A First Look at Third-Party Cyber Threat Hunting

**PI: Jamie Davis**

**Current Students: William "Trey" Maxam, USCG**

*Funding Source: Funding being sought*

### Overview

Cyber security is a huge concern for governments and private corporations alike. The average cyber intrusion costs $13 million and continues increasing. Dwell time, the amount of time an adversary is able to maintain a foothold in a compromised network undetected, is a key indicator showing organizations how capable they are of detecting adversaries once they are inside the network. Reducing dwell time significantly reduces the cost of cyber intrusions, however, according to a 2020 IBM report the average dwell time is 207 days and increasing. According to a 2019 Attivo survey (n=927), over half of the respondents said that a 100 day dwell time was either "about right" or "low". These reports highlight the importance of searching for adversary activity internal to the network boundary.

One method of detecting adversaries internal to the network boundary is a Cyber Threat Hunt (TH). Threat Hunting is "a focused and iterative approach to searching out, identifying and understanding adversaries internal to the defender's networks". Although most organizations value hunting as evidenced by widespread implementation, as a domain it is still in it's infancy with less than half of organizations utilizing a written TH process. Although theoretical threat hunt methodologies exist, the processes actually being used by TH teams are not well documented.

As a first step in understanding this important cybersecurity process, we are conducting interviews. We are interviewing threat hunt practitioners across 2 different government organizations to understand (1) their threat hunt process; and (2) the integration of non-expert and expert team members (a specific problem in the government context). Our analysis has the goal of understanding the TH process used by each practitioner and the system encompassing their processes.

### Representative Publications

*An Interview Study on Third-Party Cyber Threat Hunting Processes in the U.S. Department of Homeland Security*.
Maxam and **Davis**.
Proceedings of the 33rd USENIX Security Symposium (SECURITY'24) 2024.

# Eliminating Regex-based Denial of Service

**PIs: Jamie Davis, Dongyoon Lee (Stony Brook University)**
**Current Students: Berk Cakar (PhD student)**
*Funding Source: National Science Foundation (NSF)*

## Overview

Regexes are implemented with exponential worst-case time complexity. When used for input processing, slow regexes comprise a denial of service vector. Researchers have reported that thousands of major websites are vulnerable. This project is investigating how best to discover and eliminate these vulnerabilities. We conduct empirical work to measure vulnerability incidence in practice. We propose novel algorithms with provable security guarantees. We are exploring their adoption in production-grade regex engines.

## Representative Publications

*Using Selective Memoization to Defeat Regular Expression Denial of Service (ReDoS)*. **Davis**, Servant, and Lee. IEEE S&P 2021.

*Why Aren't Regular Expressions a Lingua Franca? An Empirical Study on the Re-use and Portability of Regular Expressions*. **Davis**, Michael, Coghlan, Servant, and Lee. ESEC/FSE 2019.

*Regexes are Hard: Decision-making, Difficulties, and Risks in Programming Regular Expressions*. Michael, Donohue, **Davis**, Lee, and Servant. ASE 2019.

*Testing Regex Generalizability And Its Implications: A Large-Scale Many-Language Measurement Study*. **Davis**, Moyer, Kazerouni, and Lee. ASE 2019.A Sense of Time for JavaScript and Node.js: First-Class Timeouts as a Cure for Event Handler Poisoning.
Davis, Williamson, and Lee. USENIX Security 2018.

*Testing Regex Generalizability And Its Implications: A Large-Scale Many-Language Measurement Study*. **Davis**, Moyer, Kazerouni, and Lee. ASE 2019.

*Exploiting Input Sanitization for Regex Denial of Service*. Barlas, Du, and **Davis**. ICSE 2022.

*Improving Developers' Understanding of Regex Denial of Service Tools through Anti-Patterns and Fix Strategies*. Hassan, Aamir, Lee, **Davis**, and Servant. Proceedings of the 44th IEEE Symposium on Security and Privacy (IEEE S&P) 2023.

*SoK: A Literature and Engineering Review of Regular Expression Denial of Service*. Bhuiyan, Cakar, Burmane, **Davis**, and Staicu. arXiv 2024.

Project URL: https://davisjam.github.io/

# Secured and Safe Assured Autonomy (S2A2) for Urban Air Mobility (UAM)

**PIs: Inseok Hwang, Dengfeng Sun, Shaoshuai Mou, Mahmoud Mahmoud (North Carolina Agricultural and Technical State University)**

**Current Students: Sounghwan (Eric) Hwang, Shanelle Clarke, Omanshu Thapliyal, Chan-Yuan (David) Kuo, S M Nahid Mahmud, Mohammed MynUddin (North Carolina Agricultural and Technical State University)**

*Funding Source: National Aeronautics and Space Administration (NASA)*

## Overview

Intelligent Transportation Systems traditionally consider ground-based vehicles operating on roads, streets, and highways. However, substantial technological advances in electrically powered aerial vehicles have paved the way to incorporate low-altitude, on-demand air transportation services to reduce the economic costs and commuting burdens of an ever-increasing ground traffic. This newly emerging transportation system is called the Urban Air Mobility (UAM) system. Envisioned as a highly automated, decentralized, cooperative passenger and cargo-carrying air transportation system, the UAM is expected to accommodate on-demand services, such as cargo delivery, medical service, and passenger transportation, in large cities. With such a large dependence on the interconnectivity of its physical and cyber components, the UAM is a large-scale safety-critical cyber-physical system (CPS) whose operation in highly populated cities means that its failure consequences can be disastrous. Therefore, UAM integration and operations within the National Airspace System (NAS) require developing new safe and secure decision-making and control frameworks that are resilient to cyberattacks. However, any such developed framework must allow for UAM scalability, and give guarantees on safety and cybersecurity while accommodating the inherent heterogeneity (of dynamics, physical constraints, communication network topologies, resource constraints, etc.) of the subsystems within the UAM.

To address the technical challenges toward safe and secure integration and operation, we first develop an explicit mathematical model of the UAM CPS under different cyberattacks. An explicit multi-agent system model of the UAM CPS is developed to describe the physical and logical behavior of the heterogeneous decentralized UAM system, which facilitates the identification and evaluation of different cyberattack pathways and the inherent structural vulnerabilities of the UAM CPS. Unlike random disturbances or faults that can disturb the UAM CPS, the configuration of cyberattacks (e.g., Denial-of-Service (DoS), Man-in-the-Middle (MITM), False-Data-Injection (FDI), etc.) is dependent on the attacker's knowledge and intent and can be sophisticatedly designed to degrade the safe and efficient performance of UAM operations.

With this model in hand, we use an array of mathematical tools and techniques to i) detect and identify the different types of cyberattacks that can disrupt UAM operations, ii) analyze the security and safety risks and potential consequences of these cyberattacks on UAM operations, and iii) develop resilient attack detection and mitigation algorithms for risk and contingency management to allow safe, secure, and efficient large-scale UAM operations. Such algorithms are designed to allow for theoretical and experimental analyses on the performance of the UAM CPS with respect to resiliency to cyberattacks

of varying severities. To detect and identify cyberattacks, we leverage state-of-the-art machine learning technologies to develop artificial intelligence (AI)-driven watchdog procedures that can be trusted to monitor UAM behavior in real-time for cyberattack intrusions. To mitigate the effects of these cyberattacks, we use a combination of techniques on adaptive estimation and reachable set analysis to avoid unsafe paths of malicious agents in the UAM, and self-organized topology reconfiguration for resiliency to cyberattacks. Furthermore, we benchmark, through rigorous mathematical analyses and in-house high-fidelity UAM simulations developed by our industry partners, the resilient performance of these algorithms for the operation of the UAM CPS in the presence of varying severities of cyberattacks.

All things considered, we design and develop a suite of resilient and risk-mitigating decision-making and novel control algorithms with benchmarked resiliency to varying severities of cyberattacks for theoretical and experimental validation before real-time UAM deployment.

**Representative Publications**

Jiazhen Zhou, Dawei Sun, Inseok Hwang, Dengfeng Sun, “Control Protocol Design and Analysis for Unmanned Aircraft System Traffic Management”, IEEE Transactions on Intelligent Transportation Systems, 2021.

Jiazhen Zhou, Dengfeng Sun, “Safe Link Transition for Unmanned Aircraft System Traffic Management”, AIAA SCITECH 2022 Forum.

Y. Xie, S. Mou and S. Sundaram, "Towards Resilience for Multi-Agent QD-Learning," 2021 60th IEEE Conference on Decision and Control (CDC), Austin, Texas, 2021.

G. Clarke, O. Thapliyal, S. Hwang, and I. Hwang, “Attack-Resilient Distributed Optimization-based Control of Multi-Agent Systems with Dual Interaction Networks”, 2022 AIAA SciTech Forum: Intelligent Systems, San Diego, CA, January 2022

Richmond Asiedu Agyapong, Mahmoud Nabil, Abdul-Rauf Nuhu, Mushahid I. Rasul and Abdollah Homaifar, “Efficient Detection of GPS Spoofing Attacks on Unmanned Aerial Vehicles Using Deep Learning”, 2021 IEEE Symposium Series on Computational Intelligence (SSCI)

O. Thapliyal and I. Hwang, “Learning based Cyberattack Design and Defense for Supervisory Control Systems”, 2021 European Control Conference, Rotterdam, Netherland, June 29- July 2, 2021

S. Hwang, M. Cho, S. Kim and I. Hwang, "An LMI-Based Risk Assessment of Leader–Follower Multi-Agent System Under Stealthy Cyberattacks," in IEEE Control Systems Letters, vol. 7, pp. 2419-2424, 2023

S. Hwang, M. Cho, and I. Hwang, "Resilient Tracking Control For Leader–Follower Multi-Agent Systems Against Sinusoidal Sensor Attacks: An LMI-Based Framework," in IEEE Control Systems Letters, vol.9, pp. 1123-1128, 2025

M. Cho, S. Hwang, and I. Hwang, "Risk Assessment of Multi-Agent System Under Denial-of-Service Cyberattacks Using Reachable Set Synthesis," in 2024 American Control Conference (ACC). IEEE, 2024.

S. Kim, M. Cho, S. Hwang, and I. Hwang, "Safety-Critical Control for Nonlinear Affine System With

Robustness and Attack Recovery," in AIAA SCITECH 2025 Forum, pp. 2722.

Project URL: https://s2a2.ncat.edu/TC2.html

# Cybersecurity for Unmanned Aerial System Operations in Urban Environment

**PIs: Inseok Hwang, Dongyan Xu, James Goppert**

**Current Students: Shiraz Khan, Dawei Sun, Kartik Anand Pant, Zhanpeng Yang, Hyungsub Kim, Jefferson Kim and MinHyun Cho**

*Funding Source: TII SSRC*

## Overview

Over the past decade, advancements in sensor, computation, communication, and battery technologies have led to the growth of Unmanned Aerial System (UAS) operations in the defense and commercial sectors. In the near future, UAS operations in urban areas are expected to increase in their scope, level of autonomy and technological capabilities. Through the introduction of new technologies in the UAS domain, such as cloud computation, ad-hoc mesh networking and self-navigation capabilities, new vectors for cyberattacks are introduced as well. It is thus of imminent importance to secure these technologies against adversaries seeking to disrupt UAS operations or the infrastructures that rely on them.

In urban environments, the proximity of UASs to human resources means that any collisions or failure in their operation can have adverse consequences. Moreover, the various obstacles in urban environments pose a unique challenge for UAS cybersecurity, as they can interfere with the sensing capabilities of the UAS (for e.g., due to obstruction or reflection of GPS signals by the buildings). These vulnerabilities can be exploited by attackers to design stealthy cyberattacks, which can remain undetected over a period of time while maximizing their impact. Sophisticated cyberattackers may even coordinate their attack across multiple UASs in a mesh network, taking advantage of the complex UAS-UAS interactions to achieve their objectives. In addition to preventing such attacks through better encryption and authentication schemes, it is important to have failsafe mechanisms in place which can assure the safety of UAS operations, as well as that of the surrounding human life and property.

We address these distinctive issues by combining a wide range of mathematical tools to identify, analyze and patch the various security vulnerabilities in current UASs. The cybersecurity of existing UAS autopilot firmware is ensured using automated vulnerability-discovery tools, which can find and fix various logic bugs in the code that could be exploited by adversaries. To facilitate real-time cyberattack mitigation during operation, the firmware is updated with novel algorithms which enhance the robustness and resilience of UAS navigation and control. We also develop offboard cyberattack-detection algorithms

which can be used to exploit external redundancies (such as camera-based surveillance and UAS traffic management systems) to detect the presence of adversarial agents in UAS operations. Furthermore, a system-of-systems perspective is used to discover the emergent vulnerabilities in large-scale UAS operations, which may otherwise remain unidentified due to their complex and unpredictable nature.

To test and validate the developed algorithms on real UASs, it is necessary to recreate these vulnerabilities and cyberattack-mitigation mechanisms in a controlled test environment which closely resembles the real-world setting. To this end, a mixed-reality experimental testbed is being developed at the Purdue UAS Research and Test Facility (PURT), which is equipped with the largest indoor motion capture system in the world and has capabilities such as realistic emulation of GPS signal characteristics in urban environments. The goal of the project is to implement the developed cyberattack-mitigation algorithms in conjunction with emulation of various cyberattacks in the testbed, such that the security of UAS firmware and physical components can be experimentally assured before their deployment in urban environments.

### Representative Publications

Khan, Shiraz, Inseok Hwang, and James Goppert. "Robust State Estimation in the Presence of Stealthy Cyberattacks." *2022 American Control Conference (ACC)*. IEEE, 2022.

## Scalable and Concurrent Targeted Search for Digital Forensics

**PI: Umit Karabiyik**

**Current Students: Akif Ozer**

*Funding Source: Air Force Research Laboratory (AFRL)*

### Overview

The rapid proliferation of digital evidence presents significant challenges for forensic investigations, particularly in processing vast volumes of unstructured data, including system logs, multimedia, and proprietary file formats. Traditional forensic tools often fall short in efficiently analyzing and correlating such diverse data, leading to investigative bottlenecks. To address these limitations, we introduce the Forensics Search Tool (FOREST), a novel Forensics as a Service (FaaS) framework that leverages a microservice-based architecture to enable scalable, modular, and high-performance forensic analysis. FOREST integrates large language models (LLMs) to automate file summarization, intelligent tagging, and contextual analysis, significantly accelerating insight extraction from complex datasets. Furthermore, it employs advanced similarity detection techniques to identify related files across large-scale evidence repositories, facilitating anomaly detection and pattern recognition. To enhance investigative efficiency,

FOREST incorporates graph-based visualization, enabling forensic analysts to intuitively explore relationships between digital artifacts and reconstruct event timelines with high precision. Designed for air-gapped environments, FOREST ensures data integrity and security while supporting essential forensic functions such as keyword indexing, cryptographic hash comparison, and reverse engineering with Ghidra. By combining AI-driven automation, similarity analysis, and interactive visual analytics, FOREST represents a transformative solution for modern forensic investigations, bridging critical gaps in unstructured data processing and enhancing digital intelligence capabilities.

## LLM Assisted Vulnerability Detection

**PI: Aravind Machiry**

**Overview**
We are exploring the use of LLMs to aid in vulnerability detection.

## Automation of Runway Status Light System

**PI: John Mott**

**Current Students: Luigi Dy**

**Overview**
Researchers at Purdue University have developed a simplified runway status light system to prevent runway incursion incidents. Currently, technologies to prevent these incidents are usually limited to large airports due to the high cost of radar-based detection systems. Small airports are generally reliant on air traffic controllers or in some cases, simple “see-and-avoid" approaches. By using a combination of automatic dependent surveillance-broadcast (ADS-B) data and computer vision modeling, lights can be activated in real-time to indicate runway status to pilots, vehicle operators and pedestrians. ADS-B is already mandated for most aircraft in U.S. airspace, making it widely available to systems of this type. As a result, this system is a low-cost and easy to integrate option for airports of varying capacities to improve runway communications and reduce risks of incidents. This technology is well suited to small airports looking to reduce the risk of runway incursion through a low-cost approach to automated runway status light control.

Project URL: https://polytechnic.purdue.edu/advanced-aviation-analytics-institute-for-research

# Complex Networks and Systems Resilience Against Disruption Propagation

PI: Shimon Nof

## Overview

The modern critical infrastructure includes complex networks and systems such as computer information networks, supply chains and networks, power supply networks, and water supply networks. These complex systems consist of many subsystems and components that are highly interconnected and interdependent. This means a disruption occurring in one part of a network will propagate their impacts to the connected parts, quickly cascading out-of-control unless properly responded to, through prevention, detection, and/or response. To address this challenge of complex networks and systems resilience against disruption, the following three converging lines of researches are conducted: (1) Supply Network Resilience; (2) Cyber-Physical Systems Security; (3) Collaborative Response to Disruption Propagation. Line (1) develops teaming and collaboration protocols to improve resilience in supply networks. Line (2) develops dynamic response collaboration protocols to tackle disruption propagation in cyber-physical systems. Line (3) unifies the different problem contexts and domains into a framework to allow analogical reasoning and knowledge sharing.

## Representative Publications

Nguyen, Win PV, and Shimon Y. Nof. “Collaborative Response to Disruption Propagation with Established Lines of Collaboration (CRDP/ESLOC) in Cyber-physical Systems: Informatics for Decision Support.” Procedia Manufacturing, ICPR-25, Chicago, IL August 2019

Nguyen, Win PV, and Shimon Y. Nof. “Advancing Cyber-Physical Systems Resilience: The Effects of Evolving Disruptions.” Procedia Manufacturing, ICPR-25, Chicago, IL August 2019

Nguyen, Win PV, and Shimon Y. Nof. “Collaborative response to disruption propagation (CRDP) in cyber-physical systems and complex networks.” Decision Support Systems 117 (2019): 1-13.

Nguyen, Win PV, and Shimon Y. Nof. “Resilience Informatics for Cyber-augmented Manufacturing Networks (CMN): Centrality, Flow and Disruption, Studies in Informatics and Control”, ISSN 1220-1766, vol. 27(4), pp. 377-384, 2018.

Tkach, Itshak, Yael Edan, and Shimon Y. Nof. “Multi-sensor task allocation framework for supply networks security using task administration protocols.” International Journal of Production Research 55.18 (2017): 5202-5224.

Levalle, Rodrigo Reyes, and Shimon Y. Nof. “Resilience in supply networks: Definition, dimensions, and levels.” Annual Reviews in Control 43 (2017): 224-236.

Seok, Hyesung, Kyungdoh Kim, and Shimon Y. Nof. “Intelligent contingent multi-sourcing model for resilient supply networks.” Expert Systems with Applications 51 (2016): 107-119.

Levalle, Rodrigo Reyes, and Shimon Y. Nof. “A resilience by teaming framework for collaborative supply

networks." Computers & Industrial Engineering 90 (2015): 67-85.

Levalle, Rodrigo Reyes, and Shimon Y. Nof. "Resilience by teaming in supply network formation and reconfiguration." International Journal of Production Economics 160 (2015): 80-93.

Zhong, Hao, and Shimon Y. Nof. "The dynamic lines of collaboration model: Collaborative disruption response in cyber–physical systems." Computers & Industrial Engineering 87 (2015): 370-382.

Zhong, Hao, and Shimon Y. Nof. "Dynamic Lines of Collaboration - Disruption Handling & Control." Springer, Automation, Collaboration, and E-Services (ACES) Book Series, 2020.

## Cyber-Collaborative Conflicts and Errors Prevention and Detection for Network Resilience

**PIs: Shimon Nof, Hao Zhong, Xin W. Chen, Rodrigo Reyes Levalle, Win P.V. Nguyen, Xin W. Chen, Rodrigo Reyes Levalle, Win P.V. Nguyen**

**Current Students: Vivek Sangani**

### Overview

Conflicts and errors are unavoidable disruptions in complex networks such as energy grids, supply chains, and collaborative decision support systems. We aim to design and implement real-time, AI-based algorithms for effective and efficient automated conflict and error (CE) prevention, detection, and recovery (PDR) for the resilience and security of such complex systems. The relationships of different CE play an essential role: Local CE can propagate to large-scale system damage according to CE dependencies if not handled correctly and in time. On the other hand, the structural information of the CE network can improve PDR operations. Constraint-based models are designed based on the complex network theory to define CE dependencies and provide prescriptive abstractions for real-world systems. A centralized algorithm taking advantage of network structure, a decentralized algorithm enabling parallelism with distributed PDR agents, and hybrid algorithms are designed to prevent, detect, and recover from CEs. The established and new algorithms use relationships between CE constraints to improve efficiency. Analytical studies and simulation experiments on various systems have been conducted to validate the latest algorithms and compare their performance to that of traditional algorithms. Results show that for effective PDR, new algorithms shall be used according to several performance measures: Response time, coverage ability, preventability, and damage minimization. The machine learning and alignment between algorithms and network characteristics, i.e., centralized algorithms for centralized networks and decentralized algorithms for decentralized networks, improve PDR. During the PDR operations, the collaboration between PDR agents also needs to be efficiently coordinated and optimized to minimize the potentially cascading effects of CE.

### Representative Publications

Chen, X. W. and Nof, S. Y. (2012), Agent-based error prevention algorithms. Expert Syst. Appl. 39, 1 (2012), 280-287.

Xin W. Chen, Shimon Y. Nof: Conflict and error prevention and detection in complex networks. Automatica 48(5): 770-778 (2012)

Xin W. Chen, Steven J. Landry, Shimon Y. Nof: A framework ofenroute air traffic conflict detection and resolution through complex network analysis. Computers in Industry 62(8-9): 787-794 (2011)

Landry, S. J., Chen, X. W., & Nof, S. Y. (2013). A decision support methodology for dynamic taxiway and runway conflict prevention. Decision Support Systems, 55(1), 165-174.

Zhong, H., Nof, S. Y., Filip, F. G. (2014) Dynamic Lines of Collaboration in CPS Disruption Response, 19th IFAC World Congress, August 24-29, 2014, Cape Town, South Africa.

Chen, X.W., and Nof, S.Y., "Automating Prognostics and Prevention of Error, conflicts, and Disruptions," chapter 22 in Springer Handbook of Automation, 2nd Edition, 2023, 509-531.

Chen, X.W., and Nof, S.Y., "Automating Prognostics and Prevention of Error, conflicts, and Disruptions," chapter 22 in Springer Handbook of Automation, 2nd Edition, 2023, 509-531.

Hao Zhong, Shimon Y. Nof, Dynamic Lines of Collaboration: Disruption Handling & Control, 2020, Springer ACES Series, Vol. 6

Rodrigo Reyes Levalle, Resilience by Teaming in Supply Chains and Networks, 2018, Springer ACES Series, Vol. 5

Reyes Levalle, R., and Nof, S.Y. "Resilience by Teaming in Supply Network Formation and Re-Configuration." Int. J. Production Economics 160, 2015, 80–93.

Reyes Levalle, R. and Nof, S.Y., "Resilience in supply networks: Definition, dimensions, and levels," Annual Reviews in Control, 43, 2017, 224-236.

Nguyen, W.P.V., and Nof, S.Y. Resilience Informatics for Cyber-augmented Manufacturing Networks (CMN): Centrality, Flow, and Disruption," Studies in Informatics and Control, 27(4), 2018, 377-384.

Nguyen, W.P.V., and Nof, S.Y. Collaborative response to disruption propagation (CRDP) in cyber-physical systems and complex networks, Decision Support Systems, 117, 2019, 1-13.

Nguyen, W.P.V., and Nof, S.Y. "Strategic lines of collaboration in response to disruption propagation (CRDP) through cyber-physical systems." Int. J. Production Economics, 230, p. 107865, 2020.

Project URL: https://engineering.purdue.edu/~prism/prj_cepd.shtml

# Assessing the Relationship between Hacking and Various Personality Traits

**PIs: Marcus Rogers, Dr. Donald Lynam, Dr. William Graziano**

Overview
Surveys indicate that there is an increasing risk of computer intrusion, computer crime and attacks on personal and business information. Computer criminality is a serious problem that affects individuals, businesses, and our nation's security. The current study has four specific aims. First, we explore whether deviant computer behavior is part of a larger syndrome of deviance. Much research has shown that non-computer-related delinquent/criminal activities, substance use, and early/risky sexual behavior are typically seen in the same individuals and can be considered part of a larger syndrome of deviance. Second, we examine whether the personality profiles of those committing deviant computer behaviors are similar to the profiles obtained from those who engage in more general deviance. Several meta-analyses have demonstrated that interpersonal antagonism (i.e., lack of empathy, oppositionality, grandiosity, and selfishness) and problems with impulse control are the most consistent personality correlates of a variety of antisocial and deviant behavior. Our third aim is to examine a potentially unique correlate of deviant computer behavior—Asperger's syndrome. Within the past decade, questions are emerging regarding the possibility of there being a link between computer criminality and a disorder known as Asperger syndrome. Finally, our fourth objective is to further validate certain psychometric instruments for use with the "hacker" sub-culture. Based on this project's preliminary data, the authors applied for a grant to conduct a cross-cultural comparison sample.

### Representative Publications

Seigfried-Spellar, K. & Rogers, M. (2010). Psychological Assessments and Attitudes toward Deviant Computer Behavior. American Academy of Forensic Sciences Annual Meeting (Feb, Seattle, WA) – no proceedings.

# Psycholinguistic Automated Detection Tool for Criminal Insiders

**PIs: Marcus Rogers, Dr. Eric Shaw**

### Overview

The area of insider threat and abuse has received a fair amount of attention in recent years. The core characteristic of the problem stems from the trust relationship that the attacker (insider) has with the victim (business or organization). This breach of trust is further exacerbated by the fact that often persons of trust have a high level of privilege on the systems that they attack.

The current body of research is rather weak when it comes to providing any practical means of solving the problem, other than "motherhood" statements and generic suggestions for security controls. The US Secret Service/Carniege Melon study "Insider Threat Study: Illicit Cyber Activity in the Banking and Finance Sector" and the "National Infrastructure Advisory Council's Report on The Insider Threat to Critical Infrastructures : Key Recommendations" provide some interesting observations but do little to assist in developing practical solutions for risk mitigation other than basic detective control suggestions.

**Prevention, Detection and Response**

Technical controls continue to be important, especially when coping with insider attacks and unexpected failures. However, not all insider problems can be solely addressed with IT-based defenses.

The key to effectively dealing with the insider threat is to develop meaningful behavioral risk models that allow for an "early warning system" that identifies individuals who, through a combination of risk characteristics, stressors, and opportunities, are likely candidates to become criminal or dangerous insiders. This early warning system is predicated on understanding human behavior, personal history, operating environment, and intervention to prevent the individual from continuing down the critical path. The social/cognitive/behavioral component is a fertile area for research. In this project, we integrate our expertise in computer deviance, employee disgruntlement/work place safety, and automated psycholinguistic analysis of computer communications (e-mail).

### Representative Publications

Rogers, M. (2010). The Psyche of Cyber Criminals: A psycho-social perspective. In Ghosh & Turrini (Eds.) Cybercrimes: A Multidisciplinary Perspective. Heidelberg, Germany: Springer-Verlag Law Division.

## Cyber Adversary Likelihood Project

**PIs: Gene Spafford, Courtney Falk, Rick Kennell, Dongyan Xu**

*Funding Source: Sandia National Laboratories*

### Overview

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.

## Use of Deception and Misdirection in Cyber Defense

**PIs: Gene Spafford, Mike Atallah, Saurabh Bagchi, Mohammed Almeshekah (King Saud University)**

**Current Students: Douglas Rapp**

### Overview

*Funding Source: Acalvio, National Science Foundation (NSF), National Security Agency, Sandia National Laboratories*

Deception and falsehoods have long been used in security, with such oft-used techniques as decoys, false flag operations, and double agents. Use of these mechanisms in cyber security have also been around for some time, with the use of honeypots (for example). However, the mantra of "No security through

obscurity" has perhaps steered people from possibilities.

We have developed a formal classification of deception and obfuscation techniques, and study how to develop new ones for the purpose of cyber defense.

We have developed a mechanism for covert signalling doubt in veracity during remote logins. We have built a mechanism to defend against anti-foensic tools and botnets, and investigated the utility of obfuscating patches to hide vulnerabilities.

Current research is examining whether deception can be used to increase the security of OT/ICS systems.

### Representative Publications

*Covert Channels Can Be Useful! --Layering Authentication Channels to Provide Covert Communication*; by M. Almeshekah, M. Atallah and E. Spafford; in Proceedings of the 21st International Workshop on Security Protocols; F. Stajano and J. Anderson, eds.; published and \copyright in 2013 by Springer-Verlag.

*The Case of Using Negative (Deceiving) Information in Data Protection*; M. Almeshekah and E. H. Spafford; 9th International Conference on Cyber Warfare and Security (ICCWS); 2014.

*Planning and Integrating Deception into Computer Systems Defenses*; by M. Almeshekah and E. H. Spafford; in Proceedings of the New Security Paradigms Workshop (NSPW); 2014.

*Using Deceptive Information in Information Security Defenses;* by M. Almeshekah and E. H. Spafford; in International Journal of Cyber Warfare and Terrorism(IJCWT), 4 (3), 46-58, July-September 2014, IGI Global.

*Enhancing Passwords Security using Deceptive Covert Communication*, M. Almeshekah, M. Atallah and E. H. Spafford, International Conference on ICT Systems Security and Privacy Protection, IFIP SEC'15, May 26-28, 2015, Hamburg, Germany

Avery*, J., & Spafford, E. H. (2017, May). Ghost Patches: Fake Patches for Fake Vulnerabilities*. In IFIP International Conference on ICT Systems Security and Privacy Protection (pp. 399-412). Springer, Cham.

Project URL: https://www.cerias.purdue.edu/site/deception/

# Algorithmic and Graph-Theoretic Approaches to Optimal Sensor Placement in Complex Dynamical Systems

**PI: Shreyas Sundaram** **Current Students: Lintao Ye, Nathaniel Woodford, Aritra Mitra**

*Funding Source: National Science Foundation (NSF)*

## Overview

Modern technology has led to the creation of new types of sensors that provide system operators with unique abilities to monitor or estimate the state of large-scale complex systems. Once sensors are in place, state estimates can be obtained by analyzing data gathered from the deployed sensors together with mathematical models of the system. However, as systems increase in scale and complexity, the deployment of sensors for high quality state estimation remains a bottleneck in a broad spectrum of applications ranging from microprocessors to power distribution networks and societal-scale Internet-of-Things. This project supports the creation of new sensor placement (deployment) algorithms with rigorous performance guarantees. The research will produce a new understanding of the fundamental limitations and achievable performance of sensor placement algorithms, and formulate efficient placement algorithms that perform well in the presence of sensor faults and external attacks.

## Representative Publications

H. Zhang, R. Ayoub and S. Sundaram, "Sensor Selection for Kalman Filtering of Linear Dynamical Systems: Complexity, Limitations and Greedy Algorithms." Automatica, vol. 78, pp. 202 - 210, April 2017.

A. Mitra and S. Sundaram, "Distributed Functional Observers for LTI Systems." Proceedings of the IEEE Conference on Decision and Control, Melbourne, Australia, 2017.

L. Ye, S. Roy and S. Sundaram, "On the Complexity and Approximability of Optimal Sensor Selection for Kalman Filtering." Proceedings of the American Control Conference, 2018 (to appear).

A. Mitra and S. Sundaram, "Distributed Observers for LTI Systems." IEEE Transactions on Automatic Control, 2018 (to appear).

S. Roy, M. Xue and S. Sundaram, "Graph-Theoretic Analysis of Estimators for Stochastically-Driven Diffusive Network Processes." Proceedings of the American Control Conference, Milwaukee, WI, 2018 (to appear).

N. Woodford and S. Sundaram, "Sensor Selection and Attack for State Estimation of Linear Systems with Unknown Inputs." Proceedings of the IEEE Conference on Decision and Control, Miami Beach, FL, 2018 (to appear).

L. Ye, S. Roy and S. Sundaram, "A Graph-Theoretic Approach to Optimal Sensor Placement for Kalman Filtering." Proceedings of the IEEE Conference on Decision and Control, Miami Beach, FL, 2018 (to appear).

Project URL: https://engineering.purdue.edu/~sundara2/

# Algorithms for Persistent Intelligence, Surveillance, and Reconnaissance by Mobile Platforms

**PIs: Shreyas Sundaram, Saurabh Bagchi** **Current Students: Aritra Mitra, Amritha Prasad**

*Funding Source: Air Force Research Laboratory (AFRL), Sandia National Laboratories*

### Overview

This project focuses on developing scalable algorithms for enabling teams of mobile agents (e.g., UAVs) to cooperatively perform intelligence gathering, surveillance, and reconnaissance tasks. In conjunction with researchers from the Air Force Research Laboratory, we have developed algorithms to enable UAVs to locate and capture a target moving on a road network, using only information available from Unattended Ground Sensors (UGSs) placed on the road. We are currently working with researchers from Sandia National Laboratories to develop distributed algorithms to enable teams of UAVs to cooperatively monitor and estimate the state of a dynamical process (e.g., a gas spreading over a battlefield), despite the presence of misinformation or hijacked UAVs that behave maliciously. The project will also develop new algorithms for task allocation in heterogeneous teams of mobile agents.

### Representative Publications

S. Sundaram, K. Kalyanam and D. W. Casbeer, "Pursuit on a Graph under Partial Information from Sensors." *Proceedings of the American Control Conference*, Seattle, WA, 2017.

A. Mitra and S. Sundaram, "Estimating the State of a Dynamical Process using a Mobile Agent", *Proceedings of the IEEE Conference on Decision and Control*, Miami Beach, FL, 2018 (to appear).

A. Mitra, W. Abbas and S. Sundaram, "On the Impact of Trusted Nodes in Resilient Distributed State Estimation of LTI Systems", *Proceedings of the IEEE Conference on Decision and Control*, Miami Beach, FL, 2018 (to appear).

A. Mitra, J. A. Richards, S. Bagchi and S. Sundaram, "Resilient Distributed State Estimation with Mobile Agents: Overcoming Time-Varying Measurements, Communication Losses, and Byzantine Adversaries'." *Autonomous Robots*, 2019.
Project URL: https://engineering.purdue.edu/~sundara2/

# Better Static Application Security Testing

**PI: Lin Tan**

### Overview

Static application security testing (SAST) detects vulnerability warnings through static program analysis. Fixing the vulnerability warnings tremendously improves software quality. However, SAST has not been fully utilized by developers due to various reasons: difficulties in handling a large number of reported warnings, a high rate of false warnings, and lack of guidance in fixing the reported warnings.

In this project, we collaborated with security experts from a commercial SAST product and propose a set of approaches (Priv) to help developers better utilize SAST techniques. First, Priv identifies preferred fix locations for the detected vulnerability warnings, and group them based on the common fix locations. Priv also leverages visualization techniques so that developers can quickly investigate the warnings in groups and prioritize their quality-assurance effort. Second, Priv identifies actionable vulnerability warnings by removing SAST-specific false positives. Finally, Priv provides customized fix suggestions for vulnerability warnings.

Our evaluation of Priv on six web applications highlights the accuracy and effectiveness of Priv. For 75.3% of the vulnerability warnings, the preferred fix locations found by Priv are identical to the ones annotated by security experts. The visualization based on shared preferred fix locations is useful for prioritizing quality-assurance efforts. Priv reduces the rate of SAST-specific false positives significantly. Finally, Priv is able to provide fully complete and correct fix suggestions for 75.6% of the evaluated warnings. Priv is well received by security experts and some features are already integrated into industrial practice.
Representative Publications

Towards Better Utilizing Static Application Security Testing. Jinqiu Yang, Lin Tan, John Peyton, and Kristofer A Duer. In the proceedings of the International Conference on Software Engineering, Software Engineering In Practice. Acceptance Rate: 25% (30/118)

## SafeBet: Secure, Simple, and Fast Speculative Execution

**PIs: T.N. Vijaykumar, Mithuna Thottethodi  Current Students: Conor Green, Cole Nelson**

### Overview

Spectre attacks exploit microprocessor speculative execution to read and transmit forbidden data outside the attacker's trust domain and sandbox. Recent hardware schemes allow potentially-unsafe speculative accesses but prevent the secret's transmission by delaying most access-dependent instructions even in the predominantly-common, no-attack case, which incurs performance loss and hardware complexity. Instead, we propose SafeBet which allows only, and does not delay most, safe accesses, achieving both security and high performance. SafeBet is based on the key observation that speculatively accessing a destination location is safe if the location's access by the same static trust domain has been committed previously; and potentially unsafe, otherwise. We extend this observation to handle inter trust-domain code and data interactions. SafeBet employs the Speculative Memory Access Control Table (SMACT) to track non-speculative trust domain code region-destination pairs. Disallowed accesses wait until reaching commit to trigger well-known replay, with virtually no change to the pipeline. Software simulations using SpecCPU benchmarks show that SafeBet uses an 8.3-KB SMACT per core to perform within 6% on average (63% at worst) of the unsafe baseline behind which NDA-restrictive, a previous scheme of security and hardware complexity comparable to SafeBet's, lags by 83% on average.

### Representative Publications

Conor Green, Cole Nelson, Mithuna Thottethodi, & T. N. Vijaykumar. (2023). SafeBet: Secure, Simple, and Fast Speculative Execution.

Project URL: https://arxiv.org/abs/2306.07785

# Process Coloring: Information-Flow Preserving Approach to Malware Investigation

**PIs: Dongyan Xu, Eugene H. Spafford, Xuxian Jiang, George Mason University**
**Current Students:** Ryan D. Riley, Larissa A. O'Brien

## Overview

*Funding Source: Air Force Research Laboratory (AFRL) and Disruptive Technology Office (DTO) under agreement number FA8750-07-2-0041.*

To detect and investigate computer malware attacks against critical cyber infrastructures, the following capabilities are desirable: (1) raising timely alerts to trigger a malware investigation, (2) determining the break-in point of a malware incident, i.e. the vulnerable service from which the malware infiltrates the victim, and (3) identifying all contaminations inflicted by the malware during its residence in the victim. In this project, we argue that the malware break-in provenance information has not been exploited in achieving these capabilities and thus propose process coloring, a new approach that preserves malware break-in provenance information and propagates it along operating system level information flows. More specifically, process coloring assigns a "color", a unique system-wide identifier, to each remotely accessible server process. The color will be either inherited by spawned child processes or diffused transitively through process actions. Process coloring achieves three new capabilities: color-based malware warning generation, break-in point identification, and log file partitioning. The virtualization-based implementation of process coloring enables more tamper-resistant log collection, storage, and real-time monitoring. Beyond the overhead introduced by virtualization, process coloring only incurs very small additional system overhead. Experiments with real-world malware (e.g., worms and rootkits) demonstrate the advantages of processing coloring over non-provenance-preserving tools.

## Representative Publications

Xuxian Jiang, Florian Buchholz, Aaron Walters, Dongyan Xu, Yi-Min Wang, Eugene H. Spafford, "Tracing Worm Break-in and Contaminations via Process Coloring: A Provenance-Preserving Approach", to appear in IEEE Transactions on Parallel and Distributed Systems, 2008.

Xuxian Jiang, Aaron Walters, Florian Buchholz, Dongyan Xu, Yi-Min Wang, Eugene H. Spafford, "Provenance-Aware Tracing of Worm Break-in and Contaminations: A Process Coloring Approach", Proceedings of IEEE International Conference on Distributed Computing Systems (ICDCS 2006), Lisboa, Portugal, July 2006.

Project URL: http://cairo.cs.purdue.edu/projects

# Building an Intelligent, Uncertainty-Resilient Detection and Tracking Sensor Network

**PI: David Yau** **Current Students: Wenchang Zhou, Yifei Sun**

### Overview

Detection, identification, and tracking problems that arise in applications such as the protection of people and the environment against chemical, biological, radiological, nuclear, and explosive (CBNRE) plumes, can be solved by combining the modalities of sensor and cyber networks. The sensor network provides information about physical space activities, such as the locations and movements of plume sources and targets. The cyber network provides storage and computational resources to analyze and infer, based on realistic dispersion models of plumes with respect to physical phenomena (e.g., terrain and weather effects), where the plume originated, the trajectory of its movement, and the prediction of its future movement. The cyber network also makes decisions about where to task and activate sensors regarding their sensing and communication activities.

The federal SensorNet initiative has carried out initial deployments of detection, identification, and tracking sensor-cyber network (DITSCN) in Washington D.C. and Memphis Port areas. While the deployments show initial success, we have also identified three grand systems challenges: (i) the need to achieve a convergence between the physical space sensing and cyber space computation; (ii) the need to recognize uncertainties due to measurement and modeling errors and uncaptured environmental effects early in the design process; and (iii) the need to support deeply embedded operations of a DITSCN in its total environment, and accommodate a wide range of sensor modalities, computation modules, and communication modules in an open system platform.

To address the above challenges, we will carry out research in (i) network formation by sensor selection, placement, and coverage; (ii) mobile sensor scheduling and coverage; and (iii) sensor tasking protocols with management of temporal and spatial uncertainty.

# Precise Calling Context Encoding

**PI: Xiangyu Zhang**

### Overview

*Funding Source: National Science Foundation (NSF)*

Calling contexts are very important for a wide range of applications such as intrusion detection, event logging, profiling, and debugging. Most applications perform expensive stack walking to recover contexts. The resulting contexts are often explicitly represented as a sequence of call sites and hence bulky. We propose a technique to encode the current calling context of any point during an execution. In particular, an acyclic call path is encoded into one number through only integer additions. Recursive call paths are divided into acyclic subsequences and encoded independently. We leverage stack depth in a safe way to optimize encoding: if a calling context can be safely and uniquely identified by its stack depth, we do not perform encoding. We propose an algorithm to seamlessly fuse encoding and stack depth based identification. The algorithm is safe because different contexts are guaranteed to have different IDs. It also

ensures contexts can be faithfully decoded. Our experiments show that our technique incurs negligible overhead (1.89% on average). For most medium-sized programs (<100k LOC), it can encode all contexts with just one number. For large programs, we are able to encode most calling contexts to a few numbers.
Project URL: http://www.cs.purdue.edu/~wsumner/research/cc

## Causality-Driven Mitigation of Cascading Failures in Distributed Systems

**PI: Yongle Zhang**

### Overview

Cascading failure is a major cause of large-scale outages in modern cloud systems. Such failures manifest through runaway positive feedback loops, where failures amplify and replicate through the entire system. The risk of such positive feedback loops continues to escalate with the increasing cloud system complexity, a trend driven by architectures such as serverless computing and microservices that inherently feature high degrees of concurrency and interaction. Despite their critical impact, current practices rely primarily on black-box mitigation techniques such as rate limiting and circuit breaking, which often fail to address positive feedback loops originated from internal system behaviors.

In this project, we design a white-box approach to understand, detect, and mitigate positive feedback loops using causality analysis. Since positive feedback loops are fundamentally causal loops where failures cause the same type of failures, we can identify positive feedback loops by tracking causality across internal events. We propose the following research thrusts: (1) Architecture Analysis: Identifying and mitigating intrinsic positive feedback loops. We will advance the understanding of positive feedback loops by identifying those stemming from intrinsic features of distributed systems. We will apply generic causal loop prevention – tracking causality and breaking causal loops – on intrinsic positive feedback loops and explore its limits. (2) Advanced Testing: Hunting for accidental positive feedback loops. To expose accidental positive feedback loops, we will design novel testing techniques to detect causal loops by causally stitching failure propagations discovered in di!erent fault injection experiments. (3) Runtime Defense: Controlling emerging positive feedback loops. We will design runtime causal loop detection techniques using causal inference, to detect and control positive feedback loops that escape testing. (4) Production Diagnosis: Interventional debugging of runaway positive feedback loops. We will design new diagnosis techniques to enable selective and safe online debugging for runaway positive feedback loops.

# Testing and detecting software upgrade failures in data-intensive distributed systems

PI: Yongle Zhang

## Overview

In the current big data era, Internet services are often built on top of data-intensive distributed systems, such as distributed storage systems and distributed computation framework. Distributed systems have to go through software upgrade as vendors need to add new features, improve performance, and deploy patches. With the rise of continuous deployment in the industry, the frequency of distributed system software upgrade could reach thousands of deployments in a single day in a major Internet company.

Unfortunately, distributed systems could experience upgrade failures - failures happen during software upgrade. These failures often have large-scale impact as upgrade is performed on the entire sys- tem. They are typically mitigated in the production environment with canary deployment, which slowly rollout updates from a small scale to the entire cluster and downgrade if a failure is encountered. How- ever, canary deployment easily takes hours and creates a dilemma between safe and fast upgrade. In addition, many upgrade failures have persistent impact and cannot be easily resolved by downgrading.

Despite the severe consequences of upgrade failures and challenges faced by production mitigation techniques, there are no existing testing and program analysis techniques that focus on testing and analyzing the distributed system upgrade procedure systematically. This work proposes to develop such techniques optimized to detect upgrade failures in early stages through exploring the effectiveness of unique properties of the distributed system software upgrade procedure.

# Security Awareness, Education, and Training

## Deploying Cyber Emulation, Modeling, and Analysis Tools on the SOL4CE

**PI: Berkay Celik**

**Overview**
We will develop a safe and isolated virtual environment to study and test computing systems and train cyber staff to support unique enterprise computing and control system environments at Purdue University. Leveraging prior partnerships and research, we propose to optimize Sandia-developed virtualization tools on the SOL4CE (Scalable Open Laboratory for Cyber Experimentation) computer cluster at Purdue University.

We work with CERIAS (Center for Education and Research in Information Assurance and Security) faculty and students to deploy a suite of cyber emulation, modeling, and analysis tools to support Sandia's mission.

We will additionally create a parallel effort at Purdue that complements the DOE-funded effort called CYMANII (Cyber Innovation to Secure U.S. Manufacturing led by the University of Texas-San Antonio, in which Sandia and Purdue are performers. SCEPTER-ICS tools and server systems dedicated to the CYMANII effort would be significantly enhanced with an open-source version installed and maintained on the SOL4CE platform at Purdue. This additional effort would leverage current milestones and deliverables in the existing SOW, but further support the capabilities in cyber-physical experimentation platforms. At the end of the same performance period (FY22), the performers will set up and have the configurability to interface with CYMANII project objectives. This will position Purdue and Sandia for the future capability to impact factory automation and supply chain management.

## Teaching and Assessing Threat Modeling Competence in Software Courses using Systems Thinking

**PIs: Jamie Davis, Kirsten A. Davis (Engineering Education)**

**Overview**
Computing systems face diverse and substantial cybersecurity threats. To mitigate these cybersecurity threats while developing software, engineers need to be competent in the skill of threat modeling. In industry and academia, there are many frameworks for teaching threat modeling, but our analysis of these frameworks suggests that (1) these approaches tend to be focused on component-level analysis rather than educating students to reason holistically about a system's cybersecurity, and (2) there is no rubric for assessing a student's threat modeling competency. To address these concerns, we propose using systems thinking, in conjunction with popular and industry-standard threat modeling frameworks like STRIDE, in order to teaching and assessing threat modeling competency. Prior studies suggest a holistic approach like systems thinking can be suitable for understanding and mitigating cybersecurity threats. Therefore, our goal is to develop learning modules for threat modeling through the lens of systems thinking, as well as assessment rubrics to assess STRIDE threat modeling performance and assess systems

thinking performance while conducting STRIDE. At the moment, we are collaborating with the company ThreatModeler to develop an online certification in threat modeling.

### Representative Publications

*Introducing Systems Thinking as a Framework for Teaching and Assessing Threat Modeling Competency*.
Joshi, Mukherjee, Davis, and **Davis**.
Annual Conference of the American Society for Engineering Education (ASEE'24) 2024.

## Crime Scene Surveying for IoT Investigations: National Training and Technical Assistance Program

**PIs: Umit Karabiyik, Marcus Rogers**

**Current Students: Akif Ozer, Xiao Hu, Yufeng Gong**

*Funding Source: Department of Justice (DOJ)*

### Overview

The Internet of Things (IoT) is an emerging technology connecting smart devices to the Internet. As these IoT devices diversify and rapidly evolve in size, shape, and form to integrate seamlessly into households, assessing crime scenes for such devices becomes challenging. Research efforts provide detailed guidelines for investigators but often overlook that IoT data is stored not only on devices and their applications but also on cloud storage and other connected devices through wireless communication. Furthermore, there is a lack of robust and user-friendly tools specifically designed to aid in the survey of crime scenes and facilitate investigations by detecting nearby wireless devices.

This Technical Training and Assistance (TTA) program aims to address this gap by providing research-based, scientifically, legally, and operationally sound training and developing a publicly accessible smartphone application. This application will be designed to scan crime scenes by detecting beaconing signals from nearby Wi-Fi and Bluetooth devices. An example legal procedure will also be designed to secure this evidence lawfully, establish a chain of custody, introduce relevant expert witnesses and specialized knowledge testimonies, develop persuasive exhibits, and formulate a prosecutorial strategy.

In addition, Purdue and NW3C will collaborate closely to create exemplary presentations and compelling demonstration exhibits. This collaboration is crucial in clarifying the digital evidence obtained from IoT technologies and illustrating its pivotal role in investigations, legal proceedings, and trials. These joint efforts align with the LECC's objectives of enhancing awareness, expanding educational opportunities, and bolstering the capabilities of criminal justice professionals.

Moreover, face-to-face training sessions, along with live and on-demand webinars, are planned to bridge the existing educational gap for numerous attendees. These sessions will focus on the legal procedures to preserve, acquire, and exhibit digital evidence from IoT devices. This TTA program will include online databases/resources for IoT device capabilities and vulnerabilities, and related training modules executed

with all sample presentations and exhibits. All technical assistance resources will be hosted on the secure portal maintained by NW3C and currently accessible to over 130,000 U.S. criminal justice practitioners.

Purdue and NW3C have long history of successful collaboration in developing and implementing TTA programs. This extensive experience and expertise give Purdue University the confidence that, in partnership with NW3C, it has an unparalleled ability to create and implement this national, integrated, and collaborative program. This TTA is specifically designed to support criminal investigators, digital forensic examiners, and prosecutors who are encountering the challenges of IoT evidence.

## CHEESE: Cyber Human Ecosystem of Engaged Security Education

**PIs: Baijian Yang, Rajesh Kalyanam**

### Overview

This project from Purdue University and the University of Illinois at Urbana-Champaign proposes to develop a set of cybersecurity learning tools on a dynamic, publicly available, web-based learning platform aimed at bringing together a community of cybersecurity researchers, educators, practitioners and students. By presenting these learning tools exclusively through a web interface, this project aims to enable ease of use and access to security education and hands-on activities. In addition to supplementing traditional cybersecurity instruction, the broader goal is to create a cybersecurity learning ecosystem that is continually updated with emerging trends in cybersecurity research as well as recently discovered security attacks. The researchers will seek contributions in both these areas from a broad community of students and faculty at community colleges, minority-serving institutions and organizations such as Women in Cybersecurity, while incorporating reward mechanisms that encourage contribution.

The proposed Cyber Human Ecosystem of Engaged Security Education (CHEESE) platform will employ the National Data Service (NDS) Labs Workbench and container technology to host a growing collection of user-contributed publicly accessible demonstrations of diverse cybersecurity concepts. Container technology provides completely independent work environments for each user and simplifies the subsequent transfer to other platforms with minimal effort. The NDS Workbench provides a user interface to container management, development and orchestration capabilities that greatly simplify the process of contributing and launching these containerized applications while automatically scaling resources in response to increased usage. CHEESE will engage a broad community of users ranging from faculty, high school teachers, students to researchers and professionals to request, contribute and evaluate the hosted cybersecurity demonstrations. The effectiveness of CHEESE in imparting practical cybersecurity knowledge will be evaluated in several undergraduate and graduate cybersecurity courses. Various organizations such as the IEEE, Big Data Hubs and Women in Cybersecurity will be leveraged to conduct training and outreach events using CHEESE to broaden community participation beyond the core user community. CHEESE will help develop a cybersecurity literate workforce that can learn from prior cybersecurity incidents and employ the latest tools and secure coding practices to avoid common pitfalls associated with these incidents.

# SaTC-EDU: EAGER Enhancing Cybersecurity Education Through a Representational Fluency Model

**PIs: Baijian Yang, Dr. Melissa Dark, Dr. Yingjie Chen, Dr. Samuel Wagstaff**

**Current Students: Beckman Joe, Sumra Bali, Wenjie Wu, Zhenzhi Xu, Ziyang Zhang**

## Overview

Create Cybersecurity experts with not only deep technical skills, but also the capabilities to recognize and respond to complex and emergent behavior, as well as a "security mindset", which includes mastery in using abstractions and principles, assessing risk and handling uncertainty, problem-solving, and reasoning; coupled with facility in adversarial thinking.

## Representative Publications

Beckman, J.G, Bari, S.G, Chen, Y., Dark, M. J., & Yang, B. (2017). The Impacts of Representational Fluency on Cognitive Processing of Cryptography Concepts (pp. 59-67). USENIX.

Beckman, J.G, Dark, M. J., P.G, Bari, S.G, Wagstaff, S. S., Chen, Y., & Yang, B. (2017). Cognitive Processing of Cryptography Concepts: An fMRI Study. Columbus, Ohio: ASEE 2017.

Serrano Aazco, M., Magana DeLeon, A. J., Yang, B. (2016). Employing Model-Eliciting Activities in Cybersecurity Education (pp. 9). ASEE.

Project URL: https://va.tech.purdue.edu/cyberFM/

# Building an Electronic Voting Technology Inspired Interactive Teaching and Learning Framework for Cybersecurity Education

PI: Xukai Zou

## Overview

Elections are the cornerstone of democratic societies, and key to their success is citizens who vote. Voting has a unique set of security and integrity requirements. The uniqueness and ubiquity of elections and the widespread use of e-voting systems emphasize the special role that e-voting technology can play in academic cybersecurity education in both college and high school. E-voting technology involves many specific and even conflicting requirements, has rich features, and covers a large knowledge base of cryptography, system security, and network security. These make it ideal as a basis for security curriculum development. Moreover, e-voting systems involve voters' active participation and interaction, rendering them even more suited for an interactive student learning process.

The goal of this project is to develop composable educational units, called "blocks," that can be combined in various ways. Each block, or set of blocks, covers a particular topic in such a way that the blocks build upon one another. Thus, they can be used in non-security courses to teach (for example) network or software security, or to build a computer security course. Such a course may be general, covering many aspects of security using the first few blocks from each area, or a specialized course using blocks from one particular area such as software security. Interactive blocks entice and enable students to fully engage in the entire learning process and more efficiently learn to master cybersecurity knowledge and skills. Collectively, the proposed framework will effectively improve student learning outcomes.

# CERIAS-Affiliated Purdue Laboratories and Center

In addition to individual connections with faculty members in keeping with the Center for Education and Research in Information Assurance and Security's multidisciplinary philosophy, CERIAS works closely with other centers and institutes on Purdue's West Lafayette and other regional campuses.

- Adaptive Computing Systems Lab (ACSL)
- AutoMous Lab
- Center for Cybersecurity Excellence and Infrastructure Protection
- Center for Resilient Infrastructures, Systems, and Processes (CRISP)
- Center for Science of Information
- Cyber Space Security Lab (Cyber2Slab)
- Cybersecurity & Forensics Lab
- Dependable Computing Systems Lab (DCSL)
- Discovery Advancements Through Analytics (D.A.T.A.) Laboratory
- Embedded Systems lab
- Flight Dynamics and Control/Hybrid Systems Lab
- Freedom Research Lab
- High Performance Computing and Cyberinfrastructure Research Lab
- Innovation Hub for Connected and Autonomous Transportation Technologies
- Institute for Global Security and Defense Innovation (i-GSDI)
- INtegrated Smart Energy Technology lab (INSET)
- Integrated Systems Laboratory
- Integrative Data Science Initiative
- International Center for Biometric Research (ICBR)
- Internet Systems Lab
- Institute for Physical Artificial Intelligence (IPAI)
- Lab for Research in Emerging Network & Distributed Systems (FRIENDS)
- Network Administration & Security Lab
- Network Infrastructure Design & Deployment Lab
- Network Learning and Discovery Lab
- NIPA Capstone Lab
- Open Ag Technologies Systems (OATS) Group
- PRISM: Production, Robotics, & Integration Software for Manufacturing & Management
- Product Lifecycle Management Lab
- Purdue Homeland Security Institute (PHSI)
- Purdue Military Research Institute (PMRI)
- Purdue Policy Research Institute
- Regenstrief Center for Healthcare Engineering
- Secure Software Systems
- Smart Machine & Assistive Robotics Technology (SMART) Lab
- SOL4CE (Scalable Open Laboratory for Cyber Experimentation)
- Systems-of-Systems (SoSE) Lab
- VACCINE
- Video and Image Processing Laboratory (VIPER)
- Wireless Network Applied Research Lab

**CERIAS Non-Purdue Partnership Organizations**

- National Colloquium Information Systems Security Education (NCISSE)
- Software Engineering Research Center
- National White Collar Crime Center (NW3C)

# CERIAS Strategic Partnership Program

CERIAS has an external partnership program that encourages industry and government agency participation and feedback to our research. Through this program, the Center also stays up-to-date on current commercial technology and emerging cyber threats. The Partnership Program provides or partners with early access to our research, technologies and graduates.

- Boeing
- Caterpillar
- CISA
- Cisco
- COMPLiQ
- Eli Lilly and Company
- Goldilock
- Idaho National Laboratory
- Intel
- Lawrence Livermore National Lab
- Lionfish Cyber Security
- Lockheed Martin
- ManTech
- NIST
- National Security Agency (NSA)
- Peraton
- RTX
- Rolls-Royce
- Sandia National Labs

**Lilly Endowment:**
In December of 1998, the Lilly Endowment made a $4.9 Million dollar grant to Purdue University for the new Center. Chronicle of Philanthropy lists the Lilly Endowment Inc. as one of the largest U.S. based foundations, with assets totaling more than $15 Billion dollars. The gift, spread over 3 years, provided an infrastructure for the new center to support a robust research and outreach program. The Lilly Endowment is recognized as a patron for their generous support.

Center for Education and Research in Information Assurance and Security
Purdue University
101 Foundry Drive
Convergence Center
Suite 3800
West Lafayette, IN 47906-3446
(765) 494-7806 - www.cerias.purdue.edu