Lancaster University

Cyber Security is of vital importance to modern society. This is true across a wide range of different sectors and leads to a diverse set of interesting research challenges for academia to investigate. Security has been made a cornerstone of Lancaster University's long-term strategic agenda, in terms of teaching, research and engagement. Security Lancaster will shortly be announcing a large number of new academic staff posts, which is the single largest commitment that Lancaster has ever made to a specific research area.

Lancaster University has consistently seen excellent growth in successful applicants to our MSc in Cyber Security and will have further exciting announcements to our teaching offerings. Finally, our partnerships with businesses and organisation such as the Lancashire and Greater Manchester Cyber Foundry, and playing a role in the decision to locate the National Cyber Force in the North West allow us to disseminate our research to key organisations in the Cyber Security field.

All this means that our academic community is of great importance. It is important that we continue to develop and strengthen our internal collaborations, which is what LS3 aims to promote. I hope that you enjoy learning about the research performed by your peers and have an interesting day discussing, learning and developing new ideas.

- Matthew Bradbury

Welcome to

Lancaster Symposium on Systems Security 2022

I am delighted to welcome you to the Lancaster Security Symposium on Systems Security 2022 (LS3 2022). This is the first time such an event has been held at Lancaster University, and we're hopeful that it will be the first of many more to come.

Over my years within the Systems Security Group, I have learnt from my colleagues about the many different research areas that contribute to improving cyber security. To this day, I am still learning new things despite (hopefully!) finishing my PhD by the end of the upcoming academic year. To this end, my vision and my fellow organisers' vision for LS3 is to share this experience so that, hopefully, you may also discover something new.

Cyber security has become an increasingly essential part of our digital lives. The Symposium provides an opportunity for all to discuss and share the important task of solving tomorrow's problems and ensuring that our digital future is secured and in safe hands. We are excited to announce several topics of research being presented today, including Cyber-Physical Systems, Threat Intelligence, Formal Verification, Machine Learning and more.

The LS3 organising committee wishes you an exciting and thought-inducing day!

- Alex Staves

Agenda

10:00	Registration
10:20	Welcome Talk
10:30	Keynote 1, UK Government
11:00	Salman Manzoor, Effectiveness of Moving Target Defense Techniques to Disrupt Attacks in the Cloud
11:15	Anna Dyson, Physical Manifestations of Ambiguity: The Drone in Contemporary Conflict
11:30	Savvas Kastanakis, The bad and the ugly of AS-path inference: A study on the confounding factors of Internet routing modelling
11:45	Thomas Miller, Looking Back to Look Forward: Lessons Learnt from Cyber-Attacks on Industrial Control Systems
12:00	Jiajie Zhang, Dynamic Scalable Distributed Key Management
12:15	Lewis Newsham, Next Generation Cyber Threat Intelligence

LUNCH BREAK (1hr)

13:30	Keynote 2, Cedric de Vylder, Cisco
14:00	Xavier Hickman, Examining the causality of network state variables on reinforcement learning driven motion planning
14:15	Zhengxin Yu, Adaptive and Robust Federated Meta Learning Framework Against Adversaries
14:30	Ovini Gunasekera, A Fresh Look at Composition of Cyber-Physical Systems
14:45	Yiqun Chen, Analyzing and Improving Customer-side Cloud Security
15:00	Andrew Sogokon, Inductive Invariants in Continuous Systems
15:15	Igor Ivkic, A Security Cost Modelling Framework for Cyber-Physical Systems

BREAK (15m)

15:45	Locknote, Ric Derbyshire, Capula
16:15	Closing Panel, LS^3 Organising Committee

Anna Dyson
Xavier Hickman
Ovini Gunasekera
Igor Ivkic

Cyber Physical Systems

Remote Physical Systems
Autonomous Systems
Security Cost Modelling
Reinforcement Learning

Physical Manifestations of Ambiguity: The Drone in Contemporary Conflict

Anna Dyson - Presentation at 11:15am

This paper explores the interaction between remote physical technologies and the phenomenon of ambiguity in conflict. Using the drone as an exemplar, the paper argues that remote physical systems possess characteristics that heighten ambiguity in contemporary conflict. A taxonomy of the drone is presented to reveal this ambiguous nature through an analysis of the drone's intrinsic properties of Plasticity and Cyber-physicality. Further, the paper considers how ambiguity arising from these intrinsic properties is extrinsically enhanced by the drone's increasing Ubiquity.
The machinic mediation of ambiguity in conflict presents novel opportunities for actors choosing to operate below the threshold of military retaliation. Such sub-threshold actors seek to exploit ambiguity in the ways that they operate, often to leverage plausible deniability. The multifaceted ambiguity of the drone lends itself to the opacity typifying such strategies, yet this feature of drone technology remains largely unexplored in defence and security literature. Understandings of drone use in conflict already lag far behind their development, adoption, and implementation. These issues are compounded by the velocity of change and transformation occurring as drone technology continues to advance and morph to new forms and functions. Going forward, the convergence of other disruptive technologies with remote-physical systems like drones may elevate ambiguity surrounding their use to new heights. Understanding ambiguity both as a phenomenon and as a mechanism is thus central to deepening contemporary knowledge of remote physical systems used in conflict and to inform defensive strategies related to them. The paper ultimately charts a path to understanding the drone as an inherently ambiguous device and calls for much wider consideration of the exploitation of ambiguity in relation to remote technologies utilised in conflict more broadly.

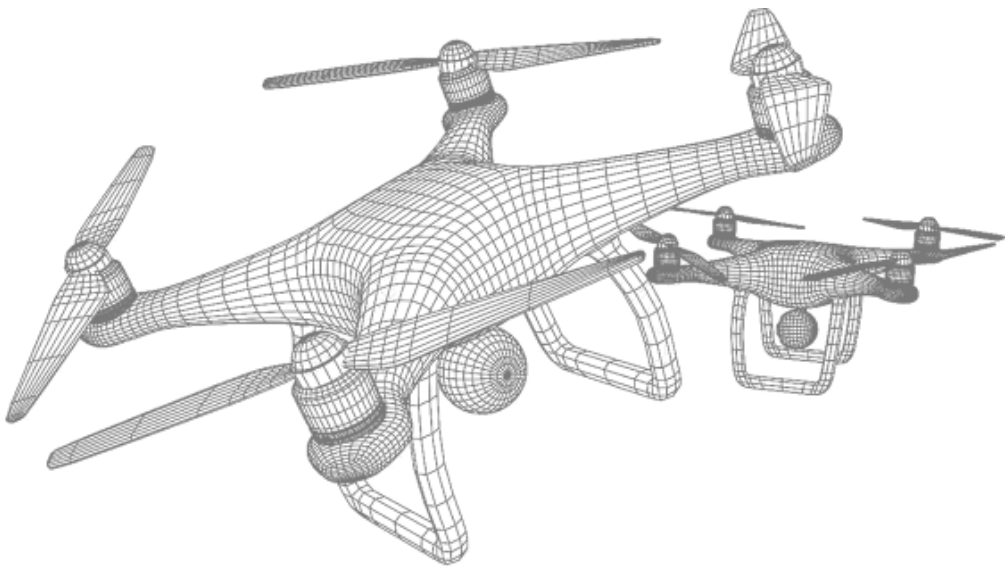

Examining the effects of channel state degradation on reinforcement learning based decision processes in autonomous agents

Xavier Hickman - Presentation at 14:00

In the modern era computational systems have a growing degree of autonomy, especially when it comes to decision making processes. The driving force behind this is the advancement of learning-based computation. Applications of this sort of technology are endless, however autonomous driving, unmanned arial vehicles (AUV) and drone navigation are a few that present some interesting challenges. Groups of decentralised fully autonomous systems like the ones previously mentioned use message passing (amongst other things) as a means of both communication and decision making. When nodes of some group of decentralised fully autonomous agents make motion decisions, they use their own individual perception of their state and their surroundings in tandem with knowledge they have on their neighbour's state and position. The data on their neighbour's state is provided via message passing and this is the aspect we are interested in. We are examining how a degrading communication medium affects the learning-based decision-making process and if the action-value function Q(st, at) is wildly mis-approximated as a result of our tests and if any specific variable of network state i.e., latency, loss, throughput has any unique affect. We are examining this relationship in the context of reinforcement learning agents using a simulated set of decentralised autonomous nodes. The nodes will have 4 degrees of freedom in a 2-dimensional Euclidean space and must navigate the space while avoiding collisions. The nodes will have varying degrees of constraints placed on their abilities with respect to both communication and sensory perception.

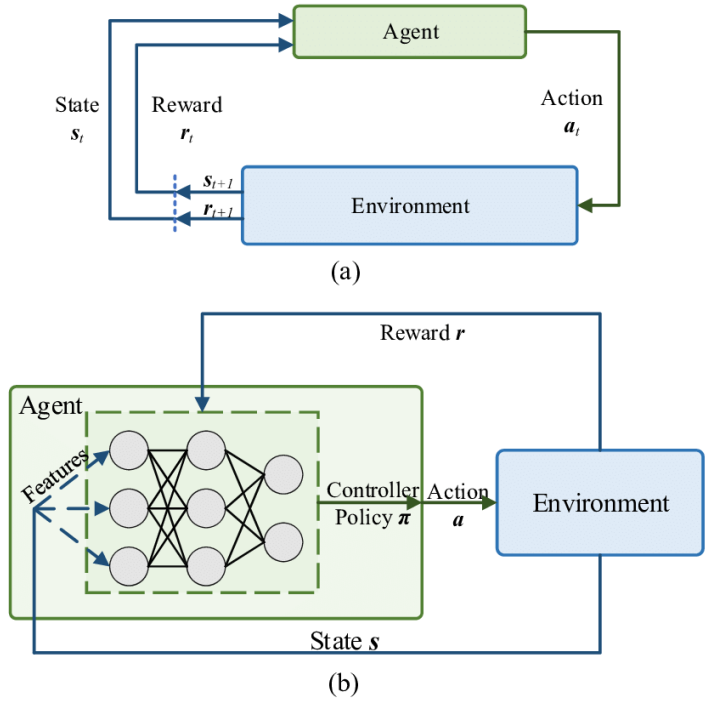

A Fresh Look at Composition of Cyber-Physical Systems

Ovini Gunasekera - Presentation at 14:30pm

A number of works have addressed the problem of modelling and reasoning about composite systems i.e. systems constituted of several individual components. Of those works that tackle the composition of Cyber-Physical Systems (CPSs), few address the concept of emergent behaviour i.e. behaviour that arises through interaction of several individual components forming a composite system. Such undetected interaction can affect the overall safety-critical operation of CPSs. This research aims to introduce a compositional specification framework which guides the specification of the interface that binds the interactions defining the composability of components. The framework attempts to formalise emergent behaviour ensuring to preserve the behaviour of individual components during composition. In order to validate the framework, a real-world problem i.e. distributed collision avoidance, will be considered as a case study from the airspace management domain.

A Security Cost Modelling Framework for Cyber-Physical Systems

Igor Ivkic - Presentation at 15:15pm

Cyber-Physical Systems (CPS) are formed through interconnected components capable of computation, communication, sensing and affecting (or changing) certain properties of the physical world. The development of these systems poses a significant challenge since they have to be designed in a way to be as secure as possible without impacting their overall performance. However, the practice of providing security always produces an additional overhead (or comes with a cost) which might lead to a compromise between performance and security. For CPS consisting of interconnected constrained devices this is even more crucial since they are often equipped with limited computational resources. To design efficient and secure CPS requires an approach for measuring and aggregating the performance of interactions including all participating components and their performed tasks. To address these challenges, we present the Security Cost Modelling Framework (SCMF) which can be used to measure, scale and aggregate the overall performance of a CPS. Unlike previous studies, our approach uses different metrics to measure the overall performance of a CPS and provides a methodology for scaling the measurement results of different units to a common Cost Unit. Furthermore, the SCMF allows extracting the overhead imposed by performing security-related tasks (or the Security Costs) from the overall performance measurements can be extracted from the overall performance measurements which allows to quantify the overhead imposed by performing security-related tasks. Using the SCMF can ultimately serve as a basis for redesigning the interactions of a CPS, so they still achieve the same overall goal, but produce less costs.

Salman Manzoor
Yiqun Chen

Cloud Security

Moving Target Defense (MTD)
Service Level Agreements (SLAs)

Effectiveness of Moving Target Defense Techniques to Disrupt Attacks in the Cloud

Salman Manzoor - Presentation at 11:00

The Moving Target Defense (MTD) approach can eliminate the asymmetric advantage that attackers have on time by changing a system's configuration dynamically to increase uncertainty and complexity for attackers. There are numerous MTDs proposed that target specific aspects of a system (e.g., IP address shuffling, dynamic OS configurations) to mitigate attacks against the component. However, deploying MTDs at different layers/components of a Cloud and assessing their effects on the overall security gains for the entire system is still very challenging due to several factors. Firstly, The Cloud is a complex system entailing physical and virtual resources. Secondly, the resources can migrate dynamically in the Cloud to satisfy user requirements. Finally, there exists a multitude of attack surfaces that an attacker can target. Therefore, we propose an MTD quantification framework that evaluates the effectiveness of MTDs for single-stage attacks. We augment the framework for multi-stage attacks. Moreover, an approach based on dynamic programming is developed to discover the optimal places in the Cloud to deploy the MTD techniques. The main contributions can be summarized as follows:

- Evaluating the needed effectiveness of MTD techniques across different layers/ abstractions of the Cloud to maximize security gains by mitigating threats.
- Quantification method to determine the effectiveness of the proposed MTD approach in disrupting both single and multi-stage attacks in the Cloud.
- Dynamic programming-based approach to explore the optimal placement for the deployment of the MTDs across the Cloud operational stack.

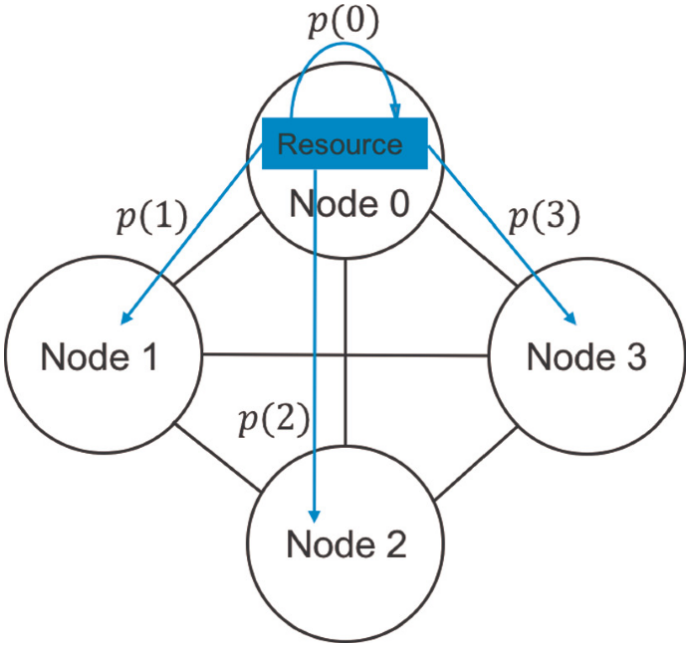

Analysing and Improving Customer-Side Cloud Security

Yiqun Chen - Presentation at 14:45

Cloud services have become popular as an effective form to outsource computational resources. While providing cost efficiency on the one side, this outsourcing also causes a certain loss of control over the computational resources, which makes security risks difficult to predict and manage. To address such concerns, security service level agreements (secSLAs) have been proposed as contracts between Cloud service providers (CSPs) and Cloud service customers (CSCs) that cover security properties of Cloud services. SecSLAs cover a variety of different security properties, ranging from the availability of encrypted communication channels for accessing Cloud resources to the timely detection and removal of vulnerabilities in the CSP's infrastructure. As previous work has shown, and as is evident for the example of timely vulnerability removal, not all of these security properties can be assessed by the CSC, which limits their utility as a contract basis. In this paper we propose a new monitoring framework for Cloud services to support the monitoring and validation of security properties on the customer side that require infrastructure-internal knowledge. To obtain the security properties to be monitored by our framework, we have manually investigated 97 different quantifiable properties in 5 standards from both industry and academia. We identified only 21 measurable properties from those standards, out of which we implement measurements for 13 representative ones and evaluated our measurements on the OpenStack platform.

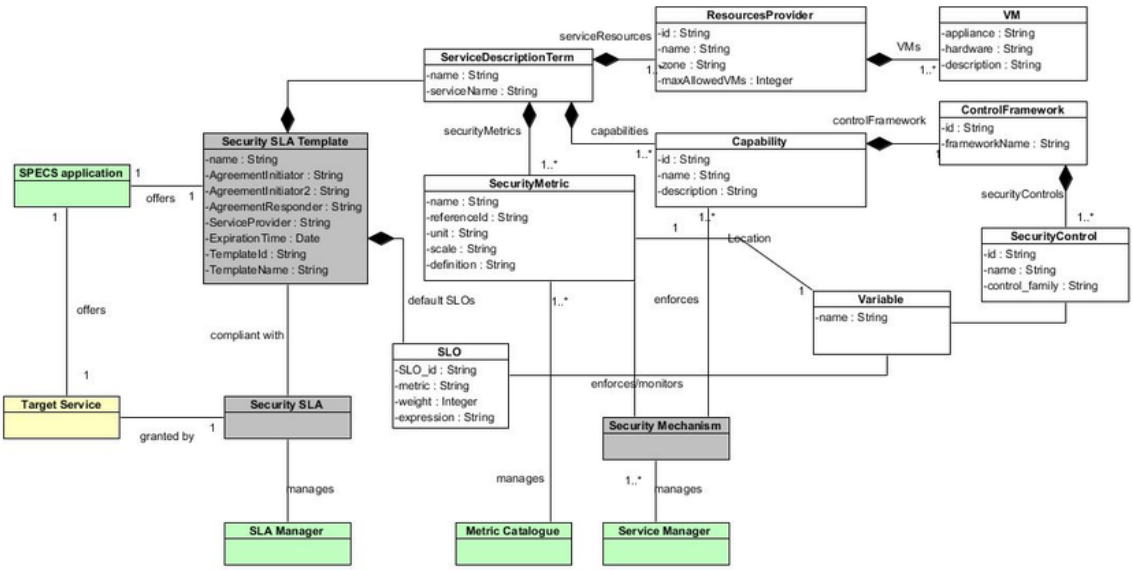

Networking

Savvas Kastanakis

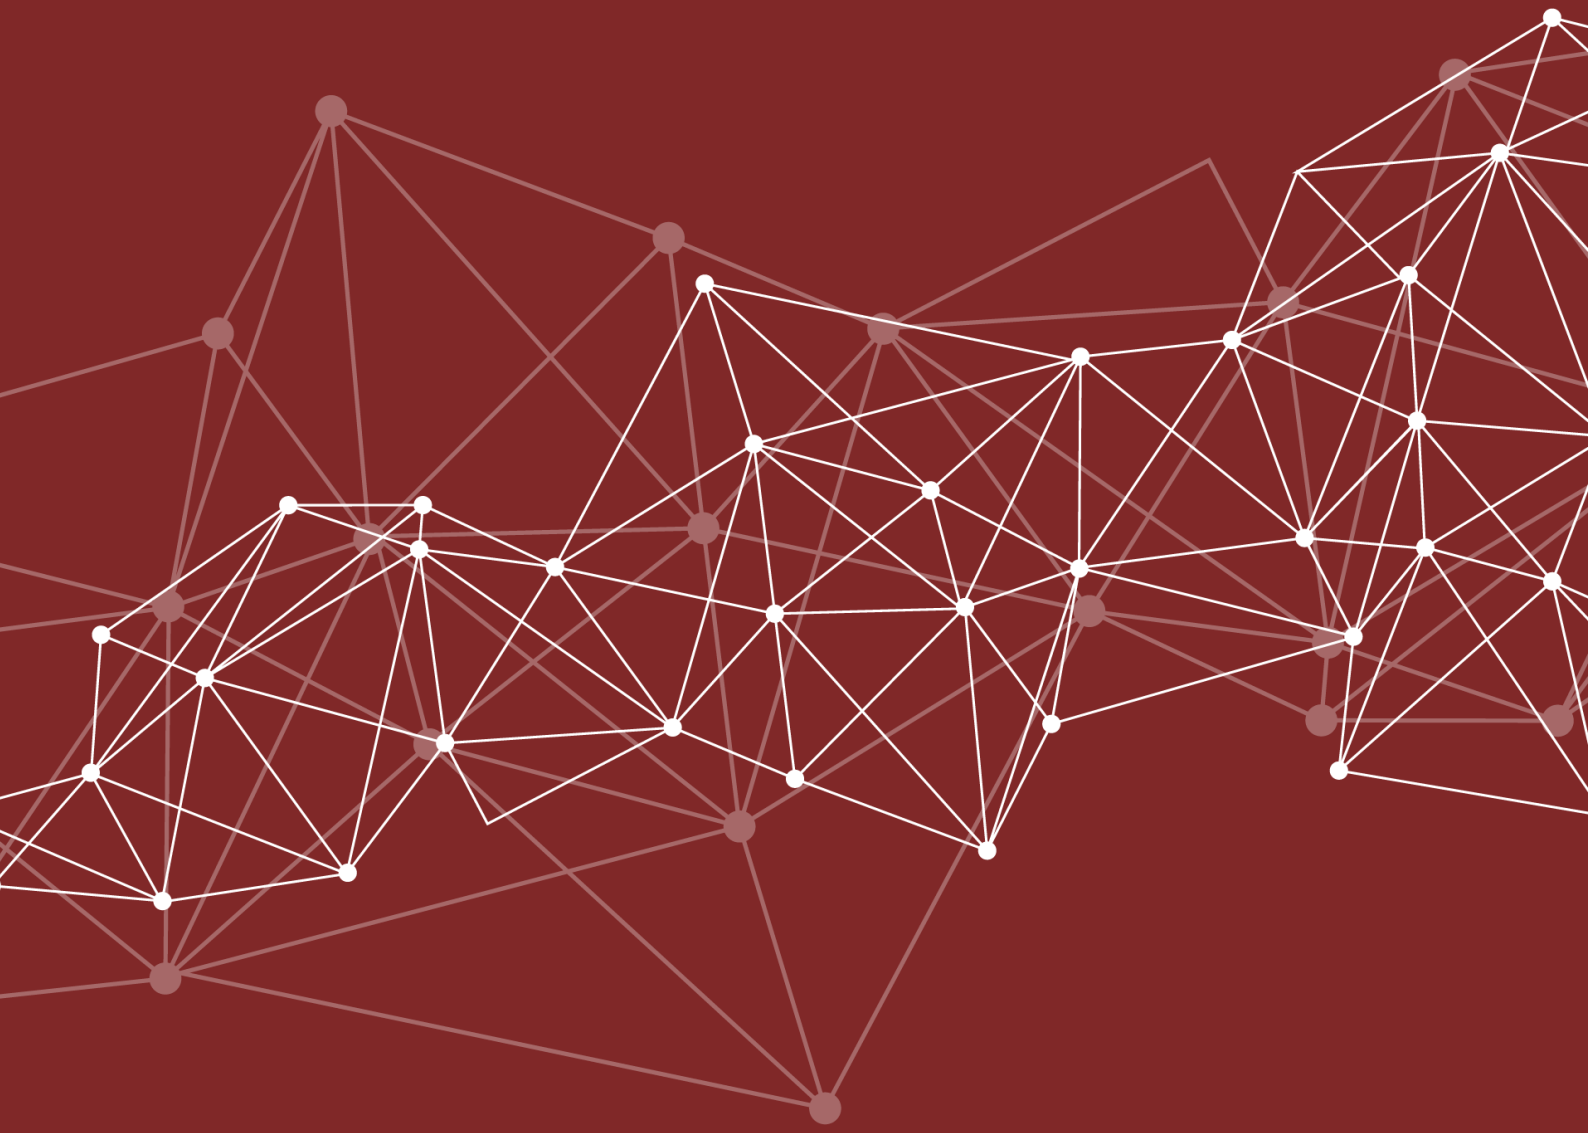

Border Gateway Protocol (BGP)
AS Routing Policies

The bad and the ugly of AS-path inference: A study on the confounding factors of Internet routing modelling

Savvas Kastanakis - Presentation at 11:30

In the broader picture, the Internet is a network of Autonomous Systems (ASes), each driven by their own business model and needs and coupled by the Border Gateway Protocol (BGP). BGP is a policy-based protocol, which enables ASes to independently define their routing policies with little or no global coordination. Hence, the forward and reverse paths between a pair of ASes may not be the same due to complex routing policies or asymmetric routing. This non-deterministic behavior of Internet routing makes the AS-level paths' inference an extremely challenging problem. Inferring the Internet's AS-level topology and paths has been of interest for a variety of reasons; network operators and researchers rely on assumptions and simulations to study new protocols, derive models to optimize network performance and evaluate network security. Even though the Internet AS-level topology has been a long-standing problem for the past two decades, an important question remains open: "which elements of Internet routing affect the AS-path inference accuracy and how much do they contribute to the error?". In this work, we are motivated by the need to provide insights on which aspect(s) of Internet routing are not modelled appropriately and how much do they contribute to the inference error. We study the Gao-Rexford model, a popular approach to infer Internet routes, aiming to identify and quantify the confounding factors that affect the inference accuracy. Our results indicate that by solving the first-hop inference problem, we can dramatically increase the exact-path score from 33.6% to 84.1% and, by further taking geography into consideration, we can refine the accuracy up to 91.1%, versus the contemporary results that are limited to 75%-85%.

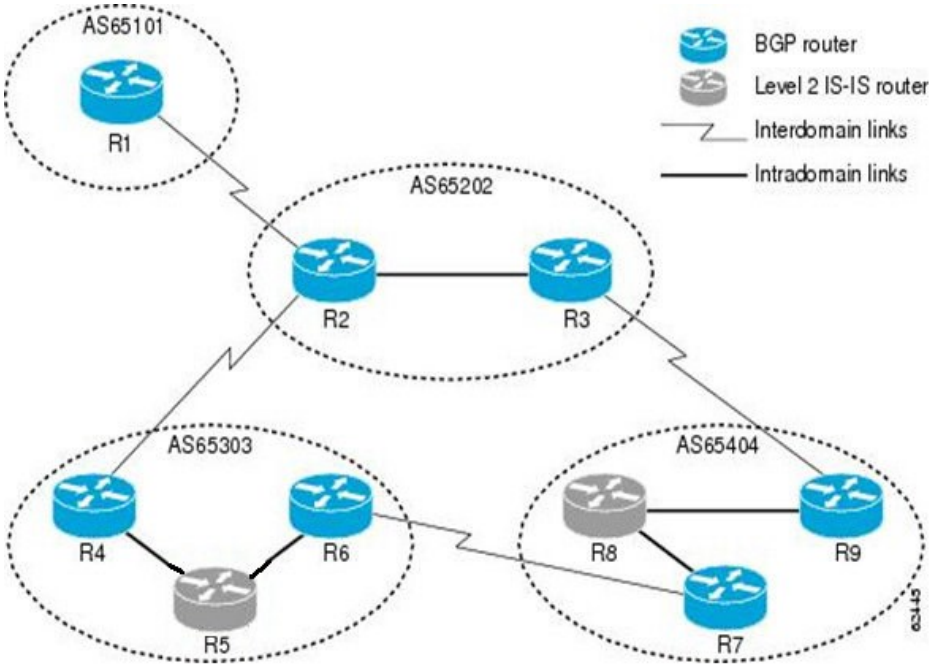

Thomas Miller
Sam Maesschalck
Alex Staves

ICS

Honeypots
Penetration Testing
Cyber-Attack Evolution

Looking Back to Look Forward: Lessons Learnt from Cyber-Attacks on Industrial Control Systems

Thomas Miller - Presentation at 11:45

Since the 1980s, we have observed a range of cyberattacks targeting Industrial Control Systems (ICS), some of which have impacted elements of critical national infrastructure (CNI). While there are access limitations on information surrounding ICS-focused cyberattacks, particularly within a CNI context, this paper provides an extensive summary of those publicly reported. By identifying and analysing previous ICS-focused cyberattacks, we document their evolution, affording cyber-security practitioners with a greater understanding of attack vectors, threat actors, impact, and targeted sectors and locations, critical to the continued development of holistic risk management strategies.

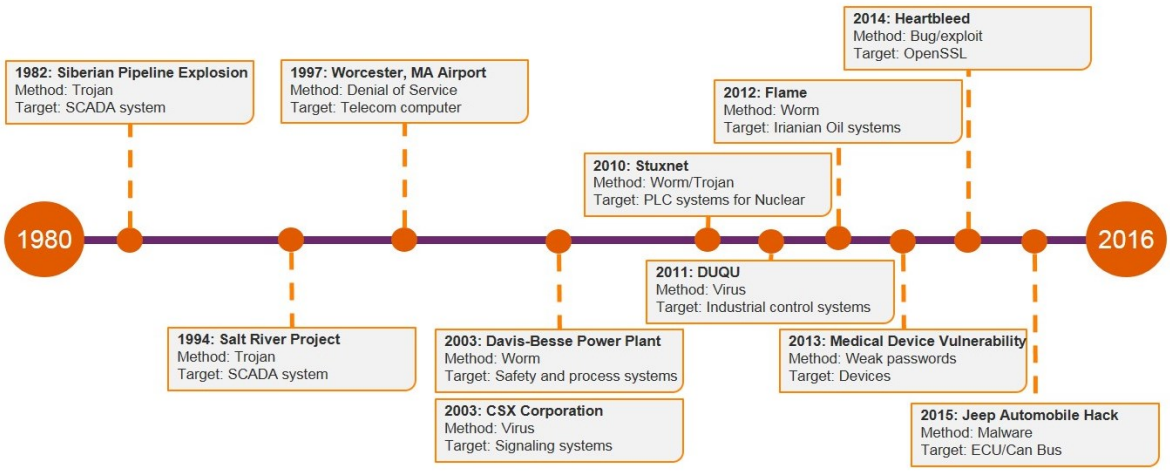

Deploying ICS Honeypots

Sam Maesschalck

The advent of Industry 4.0 and smart manufacturing has led to an increased convergence of traditional manufacturing and production technologies with IP communications. Legacy Industrial Control System (ICS) devices now interconnected via public networks are exposed to a wide range of previously unconsidered threats, which must be considered to ensure the continued safe operation of industrial processes. This paper surveys the ICS honeypot deployments in the literature to date and provides an overview of ICS-focused threat vectors, and studies how honeypots can be integrated within an organisations defensive strategy. We discuss relevant legislation, such as the UK Cyber Assessment Framework, the US NIST Framework for Improving Critical Infrastructure Cybersecurity, and associated industry-based standards and guidelines supporting operator compliance. This is used to frame a discussion on our survey of existing ICS honeypot implementations in the literature, and their role in supporting regulatory objectives. We observe that many low-interaction ICS honeypots, such as Conpot, are limited in their use and their deployment heavily influences the data obtained through them. This is largely due to the increased knowledge attackers have on how real-world ICS devices are configured and operate, vs. the configurability of simulated honeypot systems. Furthermore, we find that environments with increased interaction provide more extensive capabilities and value, due to their inherent obfuscation delivered through the use of real-world systems. We identified that the default deployment of Conpot is not enough when deploying a honeypot and explored the behaviour compared to a real PLC when conducting the reconnaissance operations. To verify red flags linked to Conpot deployments, we deploy three honeypots with a different configuration, have them scanned by Shodan and evaluate the traffic they get. Our experiments indicate that Shodan leverages CIP for ICS classification. We conclude that proper deployment of a low-interaction honeypot, such as Conpot, requires time and resources to entirely obfuscate the device and fool the attacker to a limited level. However, small changes to the default configuration does increase the performance of Conpot and results in more returning traffic. Based on these insights, we propose a novel framework towards the classification and implementation of ICS honeypots.

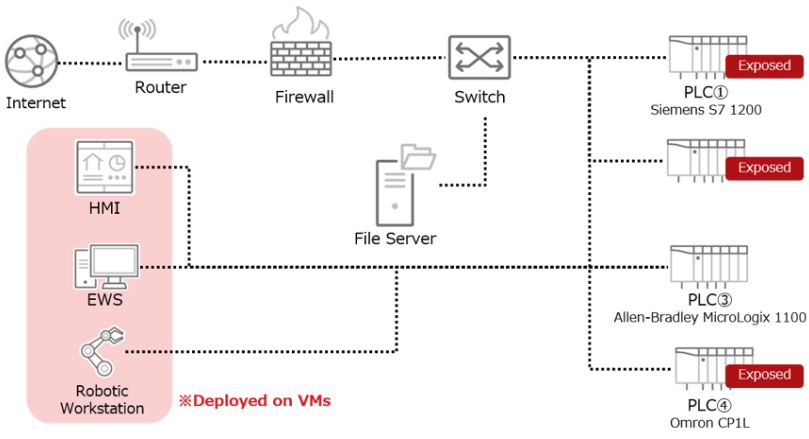

Pop! Goes the PLC: Challenges of Penetration Testing in Operational Technology Environments

Alex Staves

Assurance techniques such as adversary-centric security testing are an essential part of the risk assessment process for improving risk mitigation and response capabilities against cyber attacks. While the use of these techniques, including vulnerability assessments, penetration tests, and red team engagements, is well established within Information Technology (IT) environments, there are challenges to conducting these within Operational Technology (OT) environments, often due to the critical nature of the OT system. In this work, we provide an analysis of the technical differences between IT and OT from an asset management perspective. This analysis provides a base for identifying how these differences affect the phases of adversary-centric security tests within industrial environments. We then evaluate these findings by using adversary-centric security testing techniques on an industrial control system testbed. Results from this work demonstrate that while legacy OT is highly susceptible to disruption during adversary-centric security testing, modern OT that uses better hardware and more optimised software is significantly more resilient to tools and techniques used for security testing. Clear requirements can, therefore, be identified for ensuring appropriate adversary-centric security testing within OT environments by quantifying the risks that the tools and techniques used during such engagements present to the operational process.

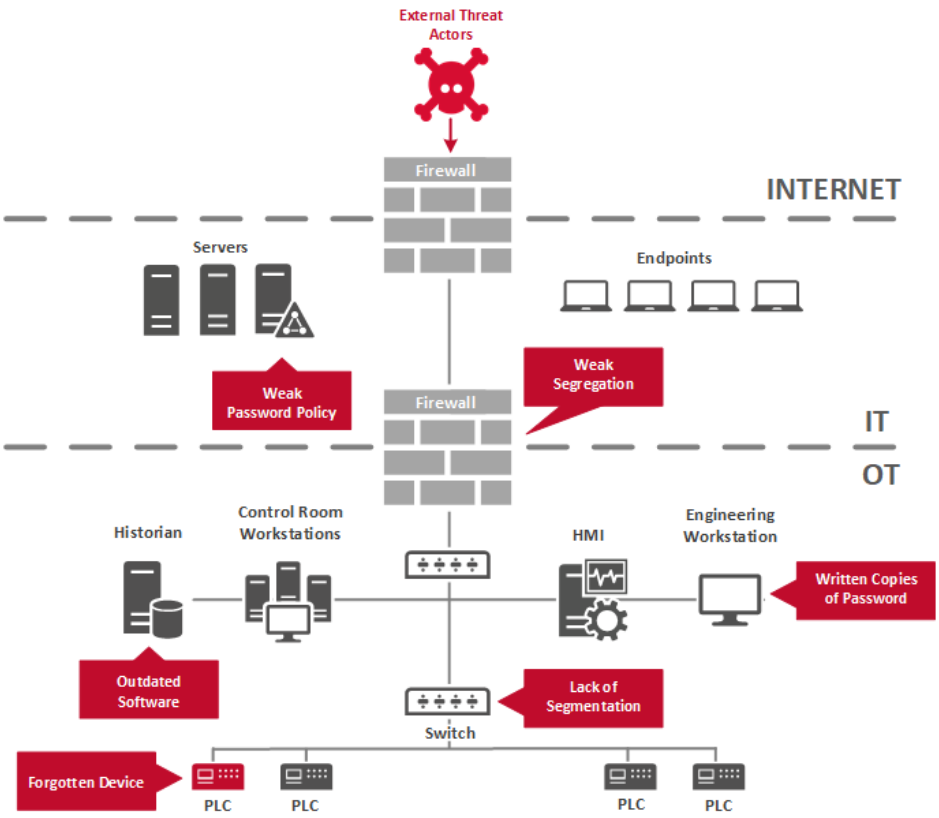

Lewis Newsham

Threat Intelligence

Proactive Intelligence
Automated Processes

Next Generation Cyber Threat Intelligence

Lewis Newsham - Presentation at 12:15

Cyber-attacks are increasingly perpetrated by organised, sophisticated and persistent entities. Even the best defenders who adopt rigorous security policies and comply with industry best practices cannot cope with the new era of next generational threats. Cyber-attack prevention thereby is now an imperative strategy for the most digital organisations, perhaps even mission-critical. The subsequent result of exponential growth in connectivity and reliance upon technology only exacerbates this issue further, whereby the divide between the cyber-offensive capability of attackers and the cyber-defensive capability of organisations becomes more asymmetric. A key avenue to resolve this asymmetry is for organisations to leverage Cyber Threat Intelligence (CTI) to better direct their cyber-defense. Effective integration of CTI can provide organisations extensive insight into the contemporary threat landscape, with added potential for realistic predictions of future threats. An inevitable by-product of virtually any CTI-led capability however, is that of data overload subsequent from excessive collection and consumption. This is due to the fact current approaches and techniques within CTI inherently require manual analysis and processing. To that end, this paper seeks to explore traditional approaches to CTI to uncover these long-standing issues in detail. An investigation into the current state-of-the-art will also be performed to shed light upon any emerging approaches and techniques that can serve in alleviating the added pressures and complexity incurred from aforementioned traditional approaches.

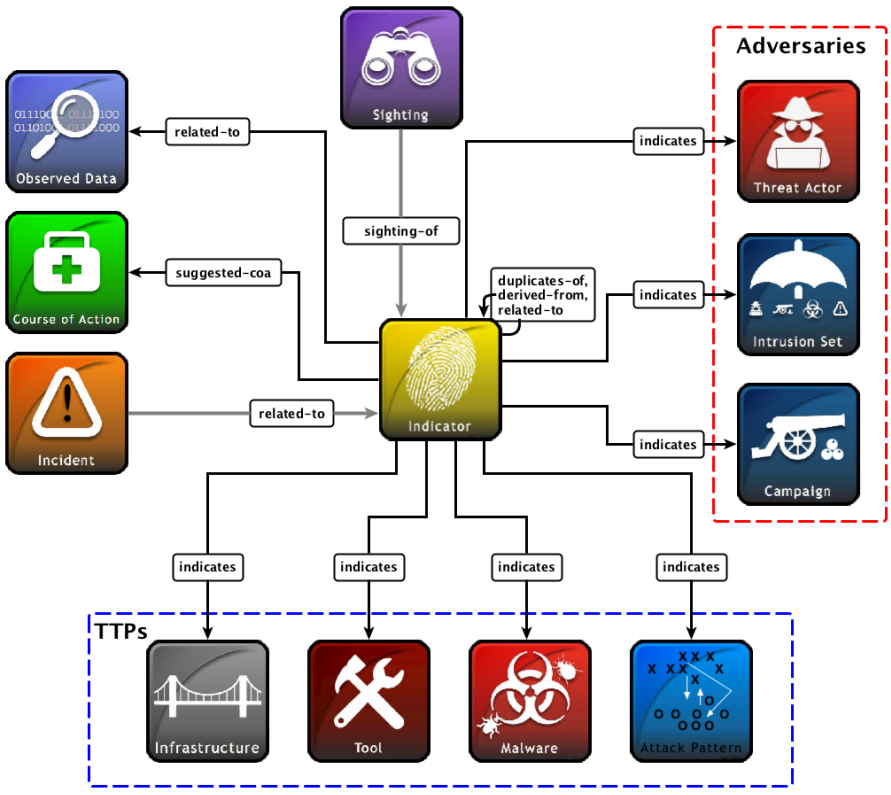

Zhengxin Yu

Machine Learning

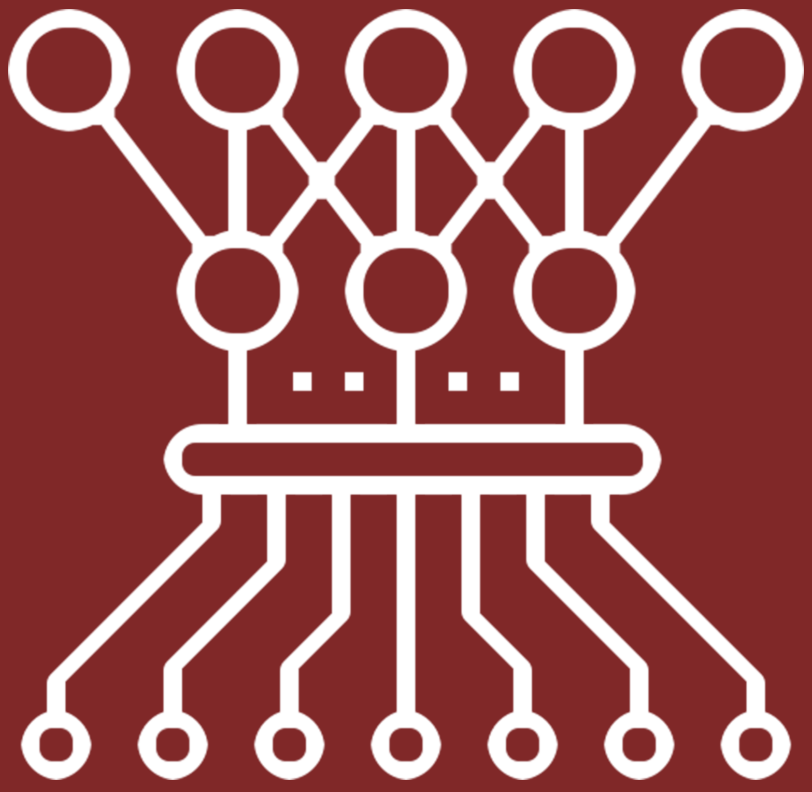

Federated Learning (FL)
Gaussian Distributions

Adaptive and Robust Federated Meta Learning Framework Against Adversaries

Zhengxin Yu - Presentation at 14:15

With the emergence of data silos and increasing privacy awareness, the traditional centralized machine learning cannot satisfy their requirements. Federated learning (FL), as a promising alternative machine learning approach, is capable of leveraging distributed personalized datasets from multiple clients to train a shared global machine learning model in a privacy-preserving manner. However, FL systems are vulnerable against clients maliciously performing poisoning attacks from uploading unreliable model updates or unintendedly uploading low-quality models, leading to degrade the FL performance. In this paper, we propose RFML: a new robust federated meta learning framework, capable of mitigating malicious model updates on non-IID data. RFML leverages a Variational AutoEncoder (VAE) to cluster clients and detect abnormal model updates and/or low-quality model updates based on their Gaussian distribution and low-dimensional embeddings. The detected malicious updates are dynamically removed from clusters. Our framework further reduces the likelihood of uploading malicious models from clients by comparing their model similarity against a calculated mean model in each cluster. Based on the similarity scores, the weighted aggregation is conducted. A small weight is given to the less model similarity. Through empirically-derived simulation of real-world scenarios, our results demonstrate that the proposed robust FL framework outperforms the conventional defense-based methods.

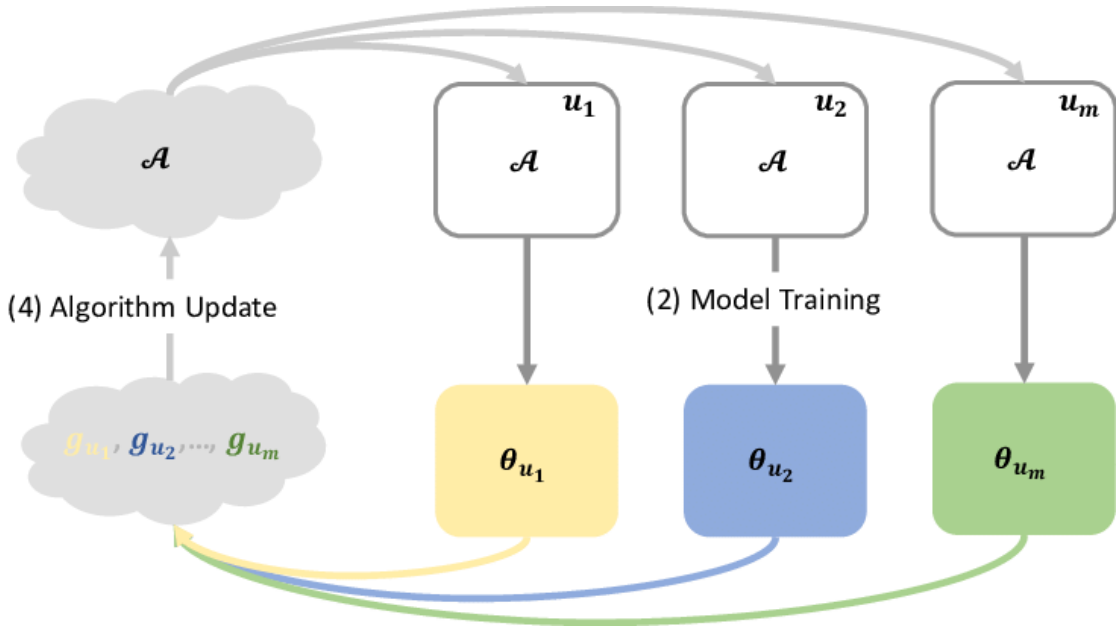

Andrew Sogokon

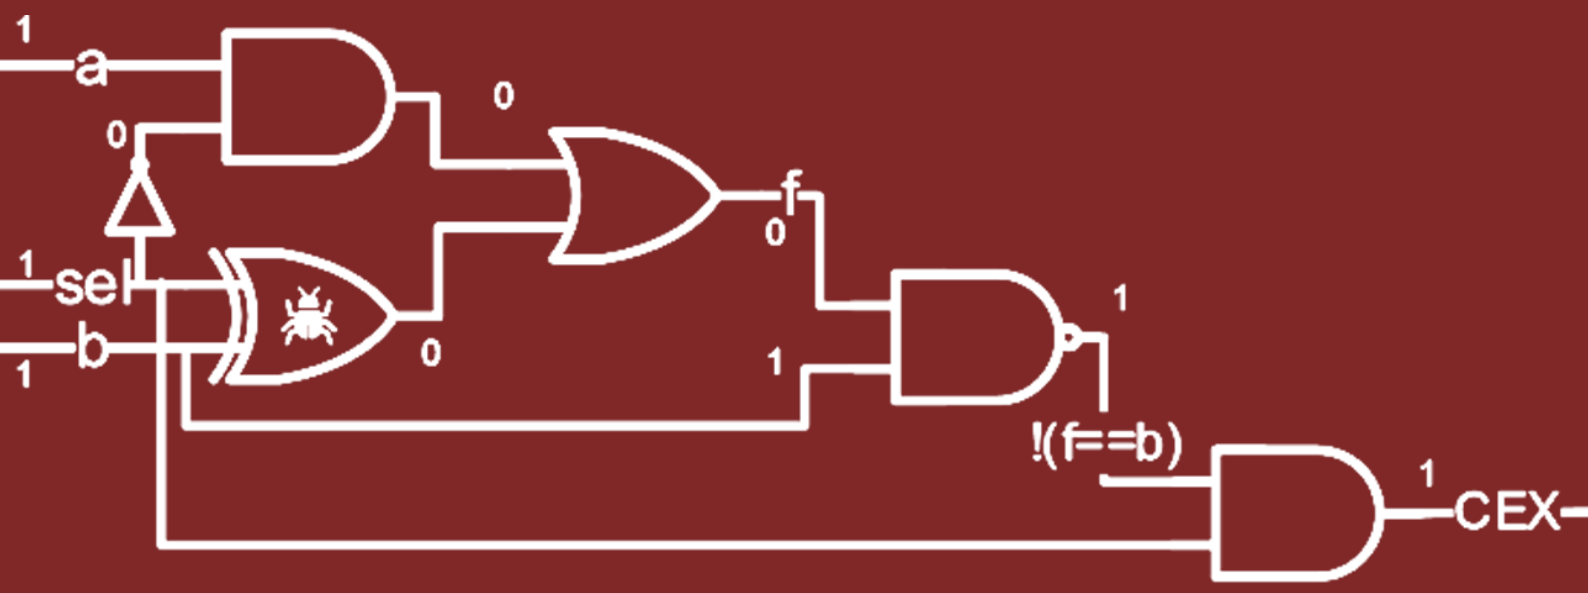

Formal Verification

Inductive Invariance
Continuous Systems

Inductive Invariants in Continuous Systems

Andrew Sogokon - Presentation at 15:00

The problem of checking inductive invariance of sets of states is of fundamental importance to formal safety verification. While the tools required for this problem in the typical discrete time setting are very straightforward and very standard in computer science, the situation is markedly different in continuous dynamical systems (i.e. dynamical systems where time does not advance in discrete time steps, but instead passes continuously, e.g. as in ordinary differential equations). This problem was considered by mathematicians since the 1940s, who developed a characterization of positively invariant sets (essentially inductive invariants); however, these results do not provide an effective way of checking whether a given set is an inductive invariant for a given continuous system. In the past decade great progress has been made by computer scientists in developing the tools for reasoning about continuous systems and powerful results have been reported which essentially solve the problem of checking inductive invariance algorithmically under some reasonable restrictions on the nature of the set and the kind of continuous system. This work will present a short overview of the current state of the art in checking inductive invariance in continuous systems and will outline some fundamental practical limitations inherent in current approaches to the problem.

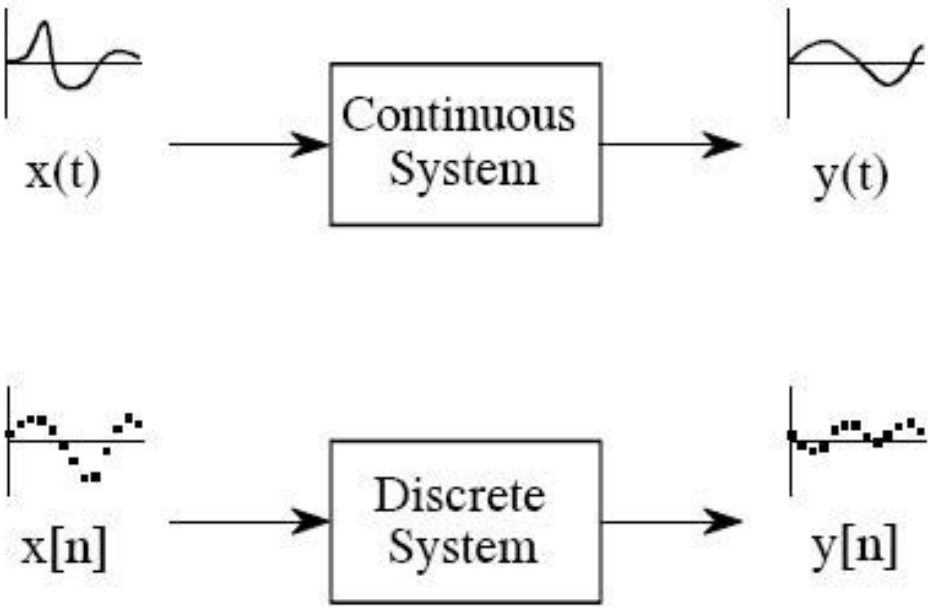

Jiajie Zhang

Encryption

Dynamic Scalable Distributed Key Management (DSKM)

Dynamic Scalable Distributed Key Management

Jiajie Zhang - Presentation at 12:00

In this paper we propose Dynamic Scalable Distributed Key Management (DSKM) scheme, which is the first DKG protocol that achieves scalability with fully dynamic participation on blockchain. Central to DSKM scheme is a Designation Procedure, and a Committee-Based Assembly line (CBA line) inspired by Algorand's player-replaceability. In Designation Procedure, Selection Committee is self-selected through cryptographic sortition. They then nominate Maintenance Committee through anonymous encryption. In CBA line, Maintenance Committees remain online to generate and maintain global distributed key pairs. Different from other works, DSKM scheme achieves scalability by employing small part of participants to generate keys instead of occupying all participants. We achieve that both committees only need to be online once in DSKM scheme. Consequently, the overall communication overhead and computational complexity of DSKM scheme is independent on the number of participants. In addition, anonymous nomination by Selection Committee further guarantees that the threshold of our DSKM scheme can scale with the number of participants.