NSF SaTC PI Meeting 2022 Breakout Session Report

Cyber-Physical Security and Privacy

Leads:
- WenZhan Song (University of Georgia)
- Alfred Chen (University of California, Irvine)

Scribes:
- Peng Liu (Pennsylvania State University)
- Sauvik Das (Georgia Institute of Technology)

Table of Contents:

1. Problem Scope and Importance

- **General problem scope**: Security and privacy of cyber-physical systems (CPSs) that (1) integrate sensing, computation, control, and networking into physical objects and infrastructure, and (2) connect with other systems, humans, and each other.
    - *How is it different from IoT?* (1) IoT is an instantiation of a CPS, and CPS is more macro-level. (2) CPS has more of a “control” aspect (e.g., closely-loop control in autonomous vehicles). (3) CPSs face more unique challenges from safety, real-time, human-in-the-loop, and privacy aspects.
- **Importance to society**:
    - Affect a wide range of *society applications*, e.g., nuclear power plants, vehicles and transportation systems, heterogeneous power grids, manufacturing systems, mobile healthcare systems, industry control systems, etc.
    - Variou levels of interplay with *human aspects*:
        - (1) Involve multiple different types of “users”, e.g., end-users (e.g., energy consumers in smart grids), operators (e.g., folks responsible for maintaining/containing/killing CPSs), deployer/integrator, etc.;

        - (2) Pervasive sensing in CPSs become intrusive to the privacy of physical-world activities;
        - (3) Unique requirement for notice & control. Specifically, adequate notice of data collection processes and control to consent/opt-out is required by privacy regulation, but CPSs generally do not have traditional UIs that would facilitate such notice and control.
- **Importance to SaTC**:
    - All the above different aspects important to society also make it important to secure and trustworthy cyberspace.
    - Aside from society impacts, from the technical security problem perspective, cyber-physical security and privacy is also important since:
        - (1) *Larger attack surface*, e.g., suffering from attack vectors from not only the cyber end, but also the physical end and the human end; the number of devices is huge and has surpassed the number of servers and desktops;
        - (2) *More diverse attack goals/targets*, e.g., attacking safety, real-time, human-in-the-loop, physical-world privacy, etc.
        - (3) *Interdependency of cyber and physical space*, e.g., some attacks may be unnoticeable in cyber or physical space alone, and data fusion and interdependent analysis are necessary;

2. Existing Research and Practices

Research

The discussions revealed various existing research lines and efforts related to cyber-physical security and privacy:

- **Analog/sensor security**
    - There is an SoK paper on this: SoK: A Minimalist Approach to Formalizing Analog Sensor Security.
- **Control security** (e.g., in ICS and drone):
    - Attack-resilient control
    - Control implementation bugs (e.g., in robotic vehicle controllers)
    - Physics-based Attack Detection (PBAD)
    - Stealthy attacks in control systems
- **AI stack security** (e.g., in autonomous driving and intelligent transportation):
    - Formal verification/analysis of AI-based autonomy
    - Semantic AI security in autonomous driving (https://sites.google.com/view/cav-sec)
    - Explanations in AI stack attacks/defenses
- **Network stack security**
    - E.g., In-Vehicle Network (IVN) security, V2X (Vehicle-To-Everything) security

- **Privacy**
  - Privacy of pervasive sensing (e.g., spying drones)
  - Location privacy for IoT/CPS one carries (or drives)
- **Real-time system aspects**
  - Real-time system availability, e.g., timing attacks on deadline violation, BFT in power grids
  - Real-time requirements for security solutions, e.g., real-time attack detection while avoiding false positives
  - Timing attacks
- **Attack/anomaly detection and defense**, for example:
  - By observing the interplay between attack detection and controller dynamic process in closed-loop control, and between cyber and physical space
  - "Real-time" security, identifying attacks as they happen by observing breaks in schedule
  - Moving target defense, e.g., by randomizing physical components and even cyber components
  - Adding honey pots/nets for attack detection
- Others that were briefly mentioned:
  - Supply chain security in CPS
  - Embedded system security

Community and Education/Training

The discussions revealed the following existing efforts on community building and education/training related to cyber-physical security and privacy:

- **Workshops**
  - Security, Privacy and Trust in IOT (co-located w/ IEEE Percom)
  - SafeThings workshop (co-located w/ IEEE S&P)
  - CPSIoTSec workshop (co-located w/ ACM CCS)
    - CPS-SPC workshop (predecessor of above)
  - AutoSec workshop (co-located w/ NDSS)
    - Record number of submissions every year for 4 consecutive years!
  - CPSS workshop (co-located wi/ ACM AsiaCCS)
    - 8 consecutive years!
- **Testbeds**
  - Testbed and digital twin for CPS security research and training (https://www.fortiphyd.com/training)
  - PASS (Platform for Autonomous driving Security and Safety): System-driven evaluation platform for autonomous driving AI security (https://sites.google.com/view/cav-sec/pass)
- **Datasets**
  - PIVOT (Platform for Innovative Use of Vehicle Open Telematics)
- **Education/training**

- AutoDriving CTF at DEFCON (https://autodrivingctf.org/)

3. Remaining Challenges and Solution Directions

The breakout session discussions were mostly devoted to discuss the important challenges in this problem space and potential solution directions. The followings are the most intensively discussed challenges:

- **Challenge: Accessibility of realistic and representative CPS security/privacy research infrastructure (testbed, dataset)**. This topic raised the most heated discussion and brainstorming. Here is a summary:
  - Different from many other cyber-only systems, CPSs usually are much more expensive end systems (e.g., autonomous vehicle) and critical infrastructure (e.g., transportation infrastructure, power grids, etc.), which makes it very hard, if not impossible (e.g., due to safety/ethical reasons), for researchers to *access a real-world CPS setup* to perform experiments.
    - Even if we can purchase one such system, CPS (e.g., ECU) reverse engineering is plagued by legal issues
  - The lack of research infrastructure is also on the *dataset* side: Today, realistic and standardized attack datasets are critically missing in CPS security. Using real infrastructure data to do research can be very risky due to the potential to affect national security.
  - Solution direction: Community-level synthetic, repeatable, sharable, standardized open testbeds for CPS security/privacy.
    - Such a testbed can also address the dataset challenge above (i.e., by generating realistic attack data from such a standardized testbed)
    - There was a consensus that NSF and/or other federal agencies should support such an effort. There were also mentions of a possible cross-agency initiative.
    - The discussions also revealed a few key challenges on building such an open testbed, for example:
      - (1) How to best approximate real-world/proprietary architectures to ensure representativeness;
      - (2) How to prevent being easily bricked, and how to handle the liability issue if such damages happened when serving for the community;
      - (3) How to simulate adversaries accurately, which should ideally be based on prior knowledge of real-world adversaries;
      - (4) How to achieve long-term maintenance.
    - It was mentioned that simulated environments (e.g., digital twin, metaverse, mixed reality) might be able to address some of the above challenges.
      - https://www.iiconsortium.org/test-beds/

- **Challenge: Creating a "science" of CPS security**. CPS systems are very variable: vehicles, power-grids, manufacturing, healthcare, IoT, etc. Can we construct a "science" of CPS security that affords a unified approach? Also there is a CPS-human security angle that is worth considering new sciences. There were comments that this is related to the Science of Security (SoS) community and CPS VO. There were suggestions that we need a CPS-Sec VO.
    - Solution direction: It was discussed that we can start by finding the basic elements of CPS security, e.g., beyond the CIA triad, such as addition of resiliency, etc.
    - **Practical vs. forward-looking threat model from science perspective**. When discussing the science of CPS, there were also discussions on the threat model considerations in CPS security research and the impacts on scientific value:
        - Specifically, the security research community often expects adversaries to be immediately practical/relevant, but that cuts off a host of forward-looking / simulated threats.
        - It was pointed out that simulated attacks can still be valuable for "science" because even if the threats have not materialized in the real-world, information about theoretical threats and mitigation strategies thereof are still important for the community
        - Agreement was reached that works can be done in different resolutions; some work can be more "practical"/"applied", while others can be more forward-looking/theoretical/simulated. All resolutions can be valid science.
- **Challenge: Ecosystem, policy, and society aspects**. Due to the inter-disciplinary and society-connected nature of CPSs, many challenges were raised from ecosystem, policy, society aspects, for example:
    - How to create incentives for investing in CPS security? Standards (e.g., AAMI/UL 2800 and 2900 series, AAMI TIR57 and TIR97, WP29, ISO 21034, etc.) may serve for this popurse. It was pointed out that certain regulators, e.g., US FDA, are in a relatively unique position to enforce safety-focused security, under the broader umbrella of holistic risk management ("Cybersecurity in Medical Devices: Quality System Considerations and Content of Premarket Submissions": https://www.fda.gov/regulatory-information/search-fda-guidance-documents/cybersecurity-medical-devices-quality-system-considerations-and-content-premarket-submissions).
    - How to effectively create, implement, and enforce security regulations and policies? There are human in the loop, which poses a big challenge to balance the criteria of multiple different stakeholder
    - How to perform risk assessment and management? CPS attacks are very low probability events. NIST risk management standards are qualitative, but it is more desired to have quantitative risk analysis.
- **Challenges: CPS security/privacy education and training**. CS students often don't have sufficient CPS domain knowledge. It is very time-consuming and difficult for CPS security/privacy because students need to be versed in security and a specific domain (e.g., vehicle networks, smart grids, healthcare, etc.)

    - Solution direction: Interdisciplinary course to train students and professionals CPS domain and security knowledge

Meanwhile, the following challenges were also raised:

- **Unique challenge from data uncertainty**. There is a lot of uncertainty in CPS data, which makes standard attack detection techniques much more challenging (e.g., anomaly detection in time series data).
- **Utility vs intrusiveness**. CPS services have privacy issues: reveal personal information to get good services. Balancing the utility provided by CPSs with the intrusiveness they entail. What's the "sweet spot"?
- **Lack of threat model summary**. There were mentions that there lacks a systematic summary of attack surface / threat models pertinent to CPS. One audience suggested that there are some papers such as the cyber-physical Dolev-Yao attack model (https://link.springer.com/chapter/10.1007/978-3-319-47846-3_12)
- **Zero-trust architectures**. There are a lot of moving parts in CPS. Can we create zero-trust architectures to simplify? At what granularity are we assuming "zero-trust"? Component-level? Device-level?
- **Biosecurity**, e.g., attacks that can change medical diagnoses by manipulating bloodwork. There are other angles such as embedding digital signatures in DNA, and security of 3D printed organs. This is a new domain and there are unique challenges on tissue engineering, e.g., the lack of a testbed with representative attacks.
- **Need to study resiliency**, i.e., what happens *when*, not *if*, CPSs are compromised by an attacker? We need failsafe defaults to help mitigate threats in the case of component/device/system failure.
- **Role of forensics in CPS attack detection**, e.g., when attacking vehicles, attackers likely leave a large trail in their wake.
- **Considerations of safety**. In CPS, there is a critical importance of physical safety, not just cyber-security. Cyber-physical systems pose real risks to physical safety of people. We need more safety-focused security (e.g., protecting users and/or operators and/or bystanders from physical and psychological damage).
- **Human privacy issues** (e.g., in the context of autonomous vehicles). In CPS, trustworthy physical world privacy protection is different from cyberspace privacy. We need to pair CPS security research with SBE.
- **Lack of understanding of industry needs**. We lack a good understanding of industry constraints in designing CPS systems. How to know the security and privacy needs from the industry?
- **Lack of system perspective in CPS research**. How to build/secure the broader CPS ecosystem as a whole instead of focusing on just closed systems (e.g., vehicles)?
- **Security metrics**. When we treat the entire CPS ecosystem (devices, operators, networks, etc.) as a whole, what kind of security measures are meaningful?

4. Other Remarks

There were also a few other remarks:

- **Industry-University-Government partnership**: Due to the nature of CPS ecosystem, it is highly desired to foster partnerships among industry, universities, and government to address their security/privacy issues from a holistic perspective.
    - https://uidp.org/
    - https://cps.uga.edu/index.php/partnership/
- **Policy and regulation aspects**: How to communicate with policymakers and law enforcement in terms of CPS security and privacy issues?
- **Academia ecosystem**. As the sub-community on this research problem space, it was recognized the general high difficulty to perform experimental evaluations directly on real-world CPS systems (e.g., commercial autonomous vehicles, transportation infrastructure, power grid, etc.). It was hoped that it can be a community consensus that theoretical attacks and open testbed based evaluation are not reasons to reject papers/proposals, since different from cyber-only security research, such a setup may already be very hard to put together and represent the best effort in this problem space.