Trust-but-Verify in Cyber-Physical Systems

Kalyan Perumalla*
perumallaks@ornl.gov
Oak Ridge National Laboratory
Oak Ridge, Tennessee, USA

ABSTRACT

Cyber-physical systems can incorporate modern technological mechanisms to harden the security of their information technology (IT) elements. However, the operational technology (OT) elements at the edge are not directly addressed by IT solutions. In this talk, we visit the core problem of adding verification of trust to existing cyber-physical systems via vetting-based verification of standards and via dynamic, passive, runtime monitoring of sensor streams to identify, characterize and monitor elements beyond the conventional IT surfaces. We illustrate how recent advances including digital twins, natural language processing and machine learning are directly useful in advancing this Trust-but-Verify approach to security and resilience of cyber-physical systems.

CCS CONCEPTS

• **Computer systems organization** → **Embedded systems**; **Sensor networks**; • **Security and privacy** → **Human and societal aspects of security and privacy**; *Network security*; • **Networks** → Network reliability.

KEYWORDS

cyber-physical systems, trust, security, vetting, verification, automation, machine learning, intelligent analysis

ACM Reference Format:
Kalyan Perumalla. 2021. Trust-but-Verify in Cyber-Physical Systems. In *Proceedings of the 2021 ACM Workshop on Secure and Trustworthy Cyber-physical Systems (SAT-CPS '21), April 28, 2021, Virtual Event, USA.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3445969.3450434

1 OVERVIEW

Cyber-physical systems span a wide spectrum, from long-lived legacy systems to more modern installations. Trust is an issue that arises across the spectrum, albeit with different variants of goals and constraints. On the one end of the spectrum, legacy systems are characterized by function-based designs in which trust is an implicitly in-built concept – the operation is historically designed, implemented, and optimized in a benign stance with respect to intended use. On the other end of the spectrum, modern systems are characterized by offerings from manufacturers, vendors, and system installers – the devices and deployments use a variety of security features that offer promises of increased trust. All along this spectrum of cyber-physical systems, extending trust beyond the traditional cyber portions to the arteries that connect the physical portions to the cyber portions is a major challenge. Here, we identify a Trust-but-Verify approach that spans this spectrum in addressing trust.

1.1 Basic Approach

In the Trust-but-Verify approach, the default view is that of trusted functions, roles, and integrated operation. However, a verification process is superimposed on this default trusted view. This is in contrast with other approaches such as those that seek to establish trust *by design* in components and installations.

The benefit of Trust-but-Verify is the ability to start with existing systems and introduce into them increasing levels of trust, which makes incremental, effective, economical and educated deployments possible. In general, this approach is most beneficial when starting with low to moderate levels of vulnerability (e.g., water treatment plants), but is not applicable to highly critical systems where the risks from breach of trust are high (e.g., nuclear plants).

1.2 Analogy

To make an analogy, the security at a local high school event employs a Trust-but-Verify approach on the visitors comprising teachers, students and their parents, in contrast to the fundamentally untrusted approach with bullet proof glasses and vests deployed in a presidential event.

Similarly, in mostly-friendly environments, trust is the default state in the operational technological part of the system, over which some verification is added to account for relatively rare and unexpected developments. As trust needs to be increased, the mechanisms are incrementally enhanced, all of which are on top of the basic foundation of trusted operation. In the high school event analogy, trust is verified with enhanced identification mechanisms and admission controls, which allows for slowly increasing costs based on current needs. On the other hand, the information technological part of the system, being much more vulnerable to severe and expected problems, is usually hardened with pervasive controls for confidentiality, integrity, and non-repudiation. In the presidential event analogy, trust is introduced and maintained only a bedrock of strong building blocks from the outset, which greatly increases costs across the board from the beginning.

SAT-CPS '21, April 28, 2021, Virtual Event, USA

ACM ISBN 978-1-4503-8319-6/21/04.
https://doi.org/10.1145/3445969.3450434

2 SPECTRUM

From the legacy end of this spectrum, trust can be added, enhanced, and maintained via the Trust-but-Verify approach by incorporating a variety of monitoring, detection, and response mechanisms. For example, existing topologies and expected behaviors can be monitored in passive modes with little perturbation to the trusted operation to gain visibility into the various dynamic components such as sensor types, value streams, their expected patterns and correlations among them. This preserves the validated nature of the existing (trusted) system while also helping perform a verification process in a delinked, asynchronous, unperturbing manner.

From the modern end of the spectrum, trust focuses not on the design, components and assembly per se, but on the advertised, promised, or assumed set of trust-related features used in the deployment. At the outset, the claims of trust-generating features from multiple parties, such as manufacturers, vendors, assemblers, installers and maintainers, are accepted as valid and representative of the devices in the installation. Nevertheless, they are vetted and verified at different levels, independently of the design or installation. These levels can span the entire gamut of available information, starting from user manuals all the way up to validation of trusted operation of actual instances of devices or sub-systems.

3 COMPUTING ADVANCEMENTS

In this spectrum of legacy to modern systems, many recent technological advancements in computing can be applied, including digital twins, natural language processing (NLP), machine learning (ML), and artificial intelligence (AI). Digital twins are useful to develop executable abstractions of legacy systems that are exercised in real-time to track normal operation versus deviations. NLP is useful to establish executable versions of user manuals, which, in a way, represent implicit agreements (or, contracts) between the device vendors and the users of the devices. ML is useful to automatically generate tolopological maps and behavioral patterns with little knowledge about the deployed systems. AI is useful to steer the Trust-but-Verify process in understanding very complex, multi-objective trust problems under time constrainted environments.

4 ILLUSTRATIONS

We are seeing advancements of Trust-but-Verify being applied in cyber-physical systems with legacy devices as well as modern ones. In systems with legacy sensor devices, trust is being added by applying intelligent analysis algorithms from computing technologies to passive instrumentation. This is achieved by extracting sensor information passively on the OT part of the system to detect, identify, and correlate sensors, in order to (a) verify the validity of the trusted configuration, and (b) detect and pin-point deviations, if any, early[3, 4]. This is especially useful for specific purposes such as triage, audit, or compliance testing. In systems being retrofitted or upgraded with more sophisticated devices (including newer remote terminal units and real-time automation controllers), the Trust-but-Verify is being explored as a way to verify the effectiveness of vendor-supplied security features, based on NLP-based processing of the vendor-advertised features and generating semi-automatic vetting engines to subject the devices to runtime tests[1, 2].

ACKNOWLEDGMENTS

This research has been supported in part by the Department of Energy Cybersecurity for Energy Delivery Systems program, project number NFE-20-07981. The author wishes to acknowledge his research collaboration with Dr. Juan Lopez in uncovering some of the insights presented here.

ABOUT THE AUTHOR

KALYAN PERUMALLA is a Distinguished Research Staff Member (Band 5, 2014) at the Oak Ridge National Laboratory (ORNL, a US Department of Energy laboratory) in the Computer Science and Mathematics Division. Dr. Perumalla holds additional appointments as Joint Full Professor in the School of Industrial and Systems Engineering at the University of Tennessee, Knoxville, and as Adjunct Professor in the School of Computational Sciences and Engineering at the Georgia Institute of Technology. He also serves on the Special Interest Group Governing Board of the Association for Computing Machinery (ACM) as the elected chair for ACM Special Interest Group in Simulation (SIGSIM). Prior to his research career at ORNL since 2005, he held full-time research appointments since 1997 at the Georgia Institute of Technology. He also served as Fellow of the Institute of Advanced Study at Durham University, UK, and as member of the National Academies' Technical Advisory Boards for the U.S. Army Research Laboratory.

REFERENCES

[1] K. Ameri, H. Sharif, J. Lopez, and K. Perumalla. 2021. Smart Semi-Supervised Accumulation of Large Repositories for Industrial Control Systems Device Information. In *2021 International Conference on Cyber Warfare and Security (ICCWS)*.

[2] K. Perumalla, J. Lopez, M. Alam, O. Kotevska, M. Hempel, and H. Sharif. 2020. A Novel Vetting Approach to Cybersecurity Verification in Energy Grid Systems. In *2020 IEEE Kansas Power and Energy Conference (KPEC)*. 1–6. https://doi.org/10.1109/KPEC47870.2020.9167562

[3] K. Perumalla, S. Yoginath, and J. Lopez. 2019. Detecting Sensors and Inferring their Relations at Level-0 in Industrial Cyber-Physical Systems. In *2019 IEEE International Symposium on Technologies for Homeland Security (HST)*. 1–5. https://doi.org/10.1109/HST47167.2019.9032891

[4] S. Yoginath, V. Tansakul, S. Chinthavali, C. Taylor, J. Hambrick, P. Irminger, and K. Perumalla. 2019. On the Effectiveness of Recurrent Neural Networks for Live Modeling of Cyber-Physical Systems. In *2019 IEEE International Conference on Industrial Internet (ICII)*. 309–317. https://doi.org/10.1109/ICII.2019.00062