

A Smart City Internet for Autonomous Systems

Gregory Falco
Stanford University
Stanford, CA USA
falco@stanford.edu

Abstract—A smart city involves critical infrastructure systems that have been digitally enabled. Increasingly, many smart city cyber-physical systems are becoming automated. The extent of automation ranges from basic logic gates to sophisticated, artificial intelligence (AI) that enables fully autonomous systems. Because of modern society’s reliance on autonomous systems in smart cities, it is crucial for them to operate in a safe manner; otherwise, it is feasible for these systems to cause considerable physical harm or even death. Because smart cities could involve thousands of autonomous systems operating in concert in densely populated areas, safety assurances are required. Challenges abound to consistently manage the safety of such autonomous systems due to their disparate developers, manufacturers, operators and users. A novel network and a sample of associated network functions for autonomous systems is proposed that aims to provide a baseline of safety for autonomous systems in a smart city ecosystem. A proposed network called the Assured Autonomous Cyber-Physical Ecosystem (AACE) would be separate from the Internet, and enforces certain functions that enable safety through active networking. Each smart city could dictate the functions for their own AACE, providing a means for enforcing safety policies across disparate autonomous systems operating in the city’s jurisdiction. Such a network design sits at the margins of the end-to-end principle, which is warranted considering the safety of autonomous systems is at stake as is argued in this paper. Without a scalable safety strategy for autonomous systems as proposed, assured autonomy in smart cities will remain elusive.

Index Terms—Smart City, Assured Autonomy, Safety, Autonomous System, End-to-End Argument, Active Network, Cyber Physical System

I. INTRODUCTION

Considering the scale at which smart cities will employ autonomous systems, safety assurances will be required. This paper proposes an active network that could facilitate the safe operation of autonomous systems in a smart city ecosystem. The argument for active networking invokes an exception in the end-to-end argument citing how performance enhancements could justify functions at the network level. In this case the performance enhancement is safety, which is a priority considering it can reduce the risk of death caused by an autonomous, cyber-physical system in a densely populated smart city environment. For the sake of spurring discussion, the paper proposes sample functions that could reside at a low level of the stack in order to guarantee certain safety parameters across disparate autonomous systems. A smart city could require all autonomous systems that wish to function within the city’s jurisdiction to operate on the city’s active network. The active network could enforce safety policies and

functions customized for their smart city ecosystem across all autonomous systems.

A. Situation

As population densities have increased in cities, there has been an incremental toll on urban critical infrastructure’s performance. To alleviate this burden by enabling efficiency and cost-savings, there has been movement towards digitizing urban systems by instrumenting critical infrastructure with embedded systems and connecting these systems to the internet. The digital instrumentation and connectivity of industrial systems is referred to as the Industrial Internet of Things (IIoT). The subsequent connectivity of urban systems and services as a result of the deployment of IIoT results in the establishment of a “Smart City” [1]. An example of a smart city capability is traffic demand management. To improve the traffic flow during peak demand hours in a city, a combination of traffic sensors, traffic lights, CCTV cameras and traditional traffic police services could be leveraged to minimize congestion, thereby reducing wait times, infrastructure wear and environmental impact (e.g. air pollution).

The instrumentation of urban systems is not new, but the mechanisms for controlling these IIoT systems are changing rapidly. Originally, these IIoT systems may have employed basic ladder logic to facilitate the control of their surroundings. The cost of processors and memory on these IIoT devices has dropped drastically which has enabled many new features for the IIoT - one of which includes the ability to conduct complex processing on the edge. One use case for edge processing is using artificial intelligence (AI) to make control decisions. Such artificial intelligence can vary from enacting AI planning instructions to facilitating distributed deep learning processes.

B. Complication

Edge AI has accelerated the capabilities of autonomous, cyber-physical systems. Smart cities are increasingly employing this technology for use cases such as autonomous transportation and distributed energy generation and demand management. Because the decision-making for these systems are completed without human oversight, in real-time, there is no room for error. For example, faulty communications between a renewable energy system and a neighborhood’s microgrid potentially caused by poor system integration could result in a brown-out. Equally disconcerting is the failure of an autonomous vehicle to perform as intended, resulting in a

traffic accident that perhaps causes physical harm or death to riders.

Faults in autonomous systems as described above could be a function of multiple issues. Because of their cyber-physical nature, faults could be caused by mechanical failures. Mechanical issues are largely avoidable through proper commissioning and maintenance plans for the technology. More disconcerting are digitally caused failures such as ones relating to faulty AI logic or malicious actors via security vulnerabilities. These digital faults are problematic because they are difficult to troubleshoot, could remain undetected in the system for a prolonged period and could cause catastrophic failures in the autonomous system. Further, such digital issues could cause cascading failures across multiple autonomous systems that have logic that rely on each other - as is the case for renewable energy systems.

Perhaps the most important distinction about autonomous systems in smart cities is that because of their cyber-physical nature, a premium is placed on the physical safety of the agent. It is critical for the control of these systems to be consistent and reliable in order to facilitate safety. Without safety as a priority, these systems could ostensibly result in human harm or death due to their cyber-physical operations. This is especially important in smart cities where hundreds or even thousands of autonomous systems are operating and potentially interacting amidst densely populated urban centers. This begs the question: how do we instill safety features in every autonomous system across a given operating landscape or ecosystem? Because of the diverse field of developers, manufacturers, operators and users of such autonomous systems, these safety parameters should be instantiated in the common element across the agents - the network.

C. Proposal

The proposed network architecture called the Assured Autonomous Cyber-Physical Ecosystem (AACE) will draw from other custom, use case-specific networks that employ active networking. The ambition of AACE is to establish a baseline of safety across all autonomous systems that reside on the network. Such a baseline would enable public trust in the safety of cyber-physical, autonomous systems that operate at scale across a smart city.

II. PRIOR ART

Since the ubiquity of the Internet, there have been calls from industry groups and various governments and their constituent agencies for custom, use case specific networks that stand alone from the global Internet. Generally, these calls have been the result of specific needs that the Internet does not fully address. This has taken various forms. For example, the space industry is currently designing the “space internet” to facilitate communications across satellite systems in space [2]. Similarly, the government of Russia has set out to build its own “unplugged”, sovereign internet [3]. The ambition with such networks is that they would be developed with a particular user group in mind and interact with the global Internet in

discreet ways - if at all. Nearly all instances are a reflection of the need for greater application control within a network.

A. Network Architecture Design

The internet was designed using a set of design principles called the end-to-end arguments [4]. The end-to-end arguments dictate that “functions placed at low levels of a system may be redundant or of little value when compared with the cost of providing them at that low level” [5]. The principle continues to dictate that providing application functionality in the network allocates unnecessary strain to all applications that interact with the network - whether they make use of the application functionality or not. It also argues that because the network may not have as much information as the application, the low level subsystem cannot perform as efficiently compared with if similar requirements were in the application layer [5].

Saltzer, Reed and Clark, the authors of the end-to-end argument, permit that there are exceptions to this rule of thumb. They assert that low level functionality could become justified as a performance enhancement [5]. A major consideration is if the performance enhancement afforded by functionality at the low level is worthwhile to all the devices that reside on the network, despite additional overhead caused by the function. This question has been grappled with by scholars discussing the merits of active networking.

Bhattacharjee, Calvert and Segura describe an extension of the end-to-end argument called “active networking” which rationalizes functionality in the network. They propose a performance model to assess the benefits of functions residing in-network versus in the application [6], [7]. They argue that placing higher level functionality in the low levels of the network is not contradictory to the end-to-end principle, rather the principle merely qualifies a notional threshold of trade-offs be met before designing functions in the network. For example, there could be use cases where certain environmental information is held within the network and timely use of such information could enhance the services of the application. They proceed to propose some functions that could be of benefit across a class of systems on a network including: time and place of congestion, global patterns across the network, and the location of packet losses [6]. These classes of functions are metadata that could have direct benefit to application functionality. To serve such functions, it would be critical that all applications could receive a benefit from these functions in spite of any overhead required to serve them. The team proceeded to calculate, using their model, a cost-benefit of employing functions on the network versus the application layer for reliable multicast and application-specific congestion control. In both cases, the team determined that employing active networks outperformed the end-to-end network implementation where functions reside on the application [6].

Research in active networks peaked in the mid 1990’s amidst the DARPA Working Group on Active Networks [8]. An insight from Calvert, one of the participants of the DARPA program, was that it is difficult to quantify benefits of active

networking; yet the benefit of improved manageability and control of the network is widely cited [9]. The notion of active networking improving the control of applications is also reflected in a survey on active networking research [10]. The survey team concludes there is a major mismatch between the rate user requirements change and the status of network infrastructure, which can be interpreted as a control issue. They suggest that the future will entail protocol components instead of stacks which are highly specialized to the features of a given use case.

B. Active Networking in Action

Despite the design thinking behind active networks being discussed in the 1990s, there seems to be few networks implemented in this manner until much later. In 2012, DARPA's Mission-Oriented Resilient Clouds (MRC) program aimed to investigate how various missions have different security and performance needs and consequentially there is the ability to make different trade-offs at run-time [11]. The various projects that emerged from this DARPA program proposed functions and their trade-offs to enable the security and resilience of mission systems [12]–[14]. Each project achieved improved security by increasing control over the mission's operating environment.

Nation-states have sought to establish increased control over networks for the sake of security (and in turn national security) as is the case for Russia's RUNET [15]. The RUNET has been discussed for over a decade, but only recently has this active network been tested "successfully" where Russia claims to have disconnected itself from the Internet on December 23, 2019 [16]. As part of this, specific control features were tested such as "the possibility of intercepting subscriber traffic and revealing information about the subscriber, blocking communication services" according to a state media document published on December 5, 2019 [16]. Russia is not the first country to work towards controlling their citizens' interactions with the Internet. Both Iran and China also have versions of their own active networks to enforce control over their citizens [17].

Active networks have also been discussed for specific industry use cases in recent years. For example, the space internet, while in many ways aims to mirror a "horizontal" infrastructure akin to the global Internet, intends to have various architectural designs to serve different purposes - namely, Earth Science Enterprise, Human Exploration and Development in Space, and Space Science Enterprise [18]. The variations of architecture to serve these different purposes are to facilitate a customized level of application control for their given use case.

Over the past decade, the energy industry has also called for their own network for energy distribution. In these cases scholars have designed energy routers to facilitate the flow and delivery of energy through distributed means [19]. Control issues are cited as a barrier to the implementation of distributed generation [20]. Considering the importance of energy distribution as critical infrastructure, it is essential for operators

to have control over the applications that facilitate energy delivery.

C. Safe Cyber-Physical Systems

Control is a major theme across network architectures that present exceptions to the end-to-end principle. Saltzer, Reed and Clark posited that low level network functionality could be justified for performance enhancement, which may well include improved control. Considering the safety of autonomous, cyber-physical systems is primarily a control issue, active networking could prove useful to facilitate safety of these systems. Further, Bhattacharjee, Calvert and Segura cost/benefit calculations for active networking suggested that functions that help to orient an application in relation to its environment (such as the location of congestion) is better situated in the network than in the application [6]. Autonomous cyber-physical systems' safe use is generally reliant on the system's ability to orient itself to its surroundings, which also points towards the use of active networking to enable safety of these systems. Perhaps the most compelling reason to consider the exception to the end-to-end principle for an autonomous systems network is because low level functionality enables safety at the expense of performance, which is certainly justifiable when lives may be at stake.

Leveson, a long-established expert in safe software systems, asserts that safety issues arise because of the interactions across "unfailed" components [21]. She describes how systems cannot be modeled or analyzed for safety in isolation of other interacting systems and that testing, simulation and standard formal verification is ineffective at assessing a system's safety. While such tests may be ineffective to assess safety on individual systems, perhaps safety is better assessed at the system ecosystem or network level. This could especially be worthwhile to assess if there is in fact active networking in place to ensure the interactions across systems are safe. In the subsequent section, we propose functionality that belongs in the autonomous system's network to enable safety.

III. ARCHITECTURAL DESIGN CONSIDERATIONS

Rather than the applications themselves be designed for safety, it will be important to provide guarantees that the entire ecosystem of applications is indeed safe by embedding safety parameters in the network. Employing active networking would help to facilitate this on a specialized AACE. Such an architectural design is uniquely applicable to when there are many autonomous systems deployed at once. A smart city is an example for when a diverse group of autonomous systems will perform at scale. A rationale for separating autonomous systems from the Internet is provided below. Then, sample functions and their associated concepts for AACE are proposed that could enable the safety of autonomous systems at scale. At this time, it is unclear if these are the right functions to include in an active network nor is it clear the best mechanism for implementing these functions. Pathways to possibly doing so include software defined networks (SDNs) or network function virtualization (NFV). It is also not yet known what

the overhead cost and relative benefit for each of these sample functions may be as this is likely partially dependent on the implementation method. These sample functions are described primarily for the sake of discussion. Functions that will be discussed include: demand management, proximal insight, reliable multicast, conflict detection and controlled failure. A rationale for each function's inclusion is included below.

A. A Separate Network

There are three principal reasons why a separate network for AACE is warranted. First, air-gapping autonomous systems from the global Internet could increase the work-factor of bad actors seeking to cause harm to these systems. While not foolproof by any means as demonstrated in previous successful attacks [22], air-gapping critical functions could increase the work-factor of a would-be attacker that intends to corrupt the safety of an autonomous system. Second, certain cyber-physical systems are very sensitive to errant network activity because they are highly calibrated and/or delicate. For example, there are anecdotal reports that systems like supervisory control and data acquisition (SCADA) controllers (employed for long-range industrial control) break during penetration testing or high network activity periods. Considering autonomous systems would also be highly calibrated considering their sophistication, it is feasible for errant network activity over the Internet to cause disruptions, thereby compromising system safety. Finally, establishing a new and separate network for autonomous systems could provide the active networking functions described below.

B. Demand Management

A given smart city will ostensibly have hundreds of thousands, if not millions of sensors and associated controllers across the urban environment. Assuming each cyber-physical, autonomous system has its own intranet or communication system like a CANbus and the IIoT components of a given autonomous system interact with other autonomous systems in a unified way, there are still likely to be thousands of autonomous systems that need to interact across an ecosystem at any time. Because of this scale of operations, it will be important for the systems to have realistic expectations of communication latency, such that it can modulate its decision-making speed accordingly.

For example, an autonomous vehicle may assume that it will receive information from a dog collar about its physical location in the middle of the road during a green light in real-time. The autonomous vehicle may allow for a 2-second reaction time window to receive new emergency location data so that it can act on this information. However, in this hypothetical situation, there are so many devices on the network at a certain time that there is a 3-second delay across the network from when the dog collar's location state is captured to when it is received by the vehicle. In such a scenario, it would be critical for the vehicle to know about this delay and account for its allowable reaction-time window accordingly (by slowing down). If the network is equipped with a demand

management function like this, autonomous systems could be better equipped to handle situations that could compromise safety due to latency.

C. Proximal Insight

Because all systems will need to interact with each other in some capacity, each should be aware of the state of the most proximal autonomous systems at any given time. More specifically, the network should be able to serve each system with information regarding the latest request of proximal systems. Such requests could provide insight to changes in physical course, data load requirements of the network or emergency signals being emitted by local systems.

For example, it could be highly valuable for a network to consistently provide all autonomous systems with data concerning other systems in distress. This could provide both a mechanism that supports autonomous system resilience facilitated by the network and a mechanism for warnings about immediate safety threats to given autonomous systems that are part of the ecosystem.

D. Reliable Multicast

Reliable delivery of data to all necessary systems will be crucial for the safety of autonomous devices. Generally, data loss is common and accepted across many networking protocols. This is especially true for point to multi-point data delivery. The safety of autonomous systems will be reliant on state communication from other autonomous systems so that any given system can orient itself in relation to the rest of the ecosystem. For this reason, it will be unacceptable for data to not reach any node across the network.

For example, if a drone plans to land in a given location, to avoid collisions, it will be critical to deliver its coordinates out to other systems. Should only a limited number of systems receive the notification of landing from the drone, there is an increase in the likelihood that a collision occurs. Therefore, systems need to be assured that all applications will reliably receive data delivered. If data is not reliably received, systems should be notified accordingly so that they can make decisions about how to handle the communication failure accordingly.

E. Conflict Detection

Considering the scale at which autonomous systems will operate in a smart city, it is inevitable that conflicting commands will be issued to various systems which may lead to a safety issue. Rather than rely on the application to catch discrepancies across commands, it would be valuable for meta-level conflict detection to occur on the network. The network could help flag conflicting information for the systems and provide validation to the command as unique, rather than waiting until the conflicting logic is processed in the application. Such a function could intend to facilitate features similar to those enabled by Lincoln Laboratory's Airborne Collision Avoidance System X (ACAS X) [23].

For example, imagine a scenario where a smart parking application dictates that the same space is available to two

different vehicles at the same time. Both vehicles then proceed to take action towards securing the open spot. It would be valuable for the network to understand the constraints of the system environment and by using planning logic conclude that it is impossible for both vehicles to park in the same spot despite both taking actions to secure the space. The network could proceed to signal to one of the vehicles that a conflict was detected and therefore the spot is no longer available. Such an action may not only spare the possibility of a collision.

F. Controlled Failure

Perhaps one of the most useful safety features of an active network for a smart city ecosystem would be to enable safe failures of the autonomous systems on the network. There could be infinite possibilities for the interaction of systems to cause unsafe conditions. Rather than having a network that can protect against all of these potential safety hazards, it would be valuable for the network to allow for the safe and controlled failure of autonomous systems. One way that a network could perhaps do this is through a distributed voting mechanism where each application is requested by the network to vote on the safety of each system present. Should all nodes return with consensus regarding a bad actor node, the network would be able to eject the “bad actor” from the network in a controlled way. A mechanism as described would help to purge the network from potentially hazardous agents and continuously prune the network. When the network determines a node should be decommissioned, it is first quarantined and allowed to complete its most recent command as long as it does not impact other nodes on the network. The application is simultaneously notified by the network that it will be cut off and should plan for a safe dismount accordingly. Equally, any autonomous system that was communicating with the bad actor would also be notified that the bad actor will be decommissioned.

An example of this could involve an off-nominal drone failure caused by a malfunctioning sensor. Because of the faulty sensor, the drone is consistently flying within a zone of protected space of other aerial devices. The drone could be voted out of the network by the other devices whose airspace was violated. The network would first notify the drone of the issue as well as notify all systems that are communicating with the drone and vice versa that the system’s network permissions will be decommissioned. Then the network will transition the drone to a quarantined sub-network that could trigger the drone’s emergency landing function. Establishing a device integrity check as described through a voting mechanism would provide increased reliability and trust across the network.

IV. DISCUSSION

Smart cities will ultimately rely on autonomous systems to facilitate critical infrastructure operations, which comes with considerable benefits. Cyber-physical, autonomous systems could yield economic efficiencies, reduce environmental impact and increase convenience for citizens. However, safety risks are sure to accompany these autonomous systems.

Safety has been addressed for autonomous systems before by establishing processes for ensuring safety in systems at the software or process level [24]–[26]. While these means for facilitating safety could be highly effective for individual systems, smart cities will realize a scale of operations of autonomous systems not previously expected. This warrants a safety enforcement mechanism that could be consistent across all systems regardless of their function, make, model or how they were developed. Employing network-level safety functions will ensure that there is a scalable way to implement a baseline level of safety across all autonomous systems. Each smart city may have a different perspective on minimum required safety features; therefore, each city could dictate which functions are included in the active network for their jurisdiction. There are likely to be cost trade-offs with employing functions at this low level, but this is justified by the universal enhancement of safety and improved control over the otherwise potentially dangerous systems.

V. CONCLUSION AND FUTURE WORK

This paper intends to provide a starting point for discussion about active networking for autonomous systems and functions that should be built into AACE. Considerable research is still needed to determine which are the most critical functions required at the network level of a smart city ecosystem to ensure safety across all autonomous devices. Then, evaluating the technical feasibility of implementing these functions on the network level is required. Also, concrete technical measures of design, verification and validation will be needed. Ultimately, a cost/benefit calculation should be completed to evaluate if the function should indeed be at a low level of the stack. Functions determined to be inappropriate to host on the network for cost or other reasons should be considered across other layers of the stack such as the computer or physical mediums (e.g. wireless).

As autonomous systems become ubiquitous and their operations scale accordingly, there will be a vital need for reliable control to ensure safety and reduce the risk of death by AI-enabled autonomous, cyber-physical systems. Not only must these systems operate in a safe way themselves, but their interactions with other systems must also be assuredly safe at scale. Therefore, establishing safety controls at the lowest common denominator across all interacting autonomous systems - the network, could provide considerable benefits given the diversity of systems, developers, manufacturers, operators and users that could otherwise lead to safety failures. Creating safety assurances by embedding certain functions in the network can take us a step closer to legions of safe, autonomous cyber-physical systems for future smart cities.

REFERENCES

- [1] G. Falco, A. Viswanathan, C. Caldera, and H. Shrobe, “A master attack methodology for an ai-based automated attack planner for smart cities,” *IEEE Access*, vol. 6, pp. 48 360–48 373, 2018.
- [2] R. Wang, T. Taleb, A. Jamalipour, and B. Sun, “Protocols for reliable data transport in space internet,” *IEEE Communications Surveys & Tutorials*, vol. 11, no. 2, pp. 21–32, 2009.

- [3] M. Ristola, "Should 'runet 2020' be taken seriously? contradictory views about cyber security between russia and the west," *Journal of information warfare*, vol. 16, no. 4, pp. 113–131, 2017.
- [4] M. S. Blumenthal and D. D. Clark, "Rethinking the design of the internet: the end-to-end arguments vs. the brave new world," *ACM Transactions on Internet Technology (TOIT)*, vol. 1, no. 1, pp. 70–109, 2001.
- [5] J. H. Saltzer, D. P. Reed, and D. D. Clark, "End-to-end arguments in system design," *ACM Transactions on Computer Systems (TOCS)*, vol. 2, no. 4, pp. 277–288, 1984.
- [6] S. Bhattacharjee, K. L. Calvert, and E. W. Zegura, "Active networking and the end-to-end argument," in *Proceedings 1997 International Conference on Network Protocols*. IEEE, 1997, pp. 220–228.
- [7] —, "An architecture for active networking," in *International Conference on High Performance Networking*. Springer, 1997, pp. 265–279.
- [8] K. Calvert, "Architectural framework for active networks," *Active Network Working Group Draft*, 1999.
- [9] K. CALVERT, "Reflections on network architecture: an active networking perspective," *ACM SIGCOMM Computer Communication Review*, vol. 36, no. 2, pp. 27–30, 2006.
- [10] D. L. Tennenhouse, J. M. Smith, W. D. Sincoskie, D. J. Wetherall, and G. J. Minden, "A survey of active network research," *IEEE communications Magazine*, vol. 35, no. 1, pp. 80–86, 1997.
- [11] H. E. Shrobe, "What if we got a do-over?" in *Proceedings of the 2012 ACM Workshop on Cloud computing security workshop*, 2012, pp. 55–56.
- [12] G. Jakobson, "Mission-centricity in cyber security: Architecting cyber attack resilient missions," in *2013 5th International Conference on Cyber Conflict (CYCON 2013)*. IEEE, 2013, pp. 1–18.
- [13] A. Newell, D. Obenshain, T. Tantillo, C. Nita-Rotaru, and Y. Amir, "Increasing network resiliency by optimally assigning diverse variants to routing nodes," *IEEE Transactions on Dependable and Secure Computing*, vol. 12, no. 6, pp. 602–614, 2014.
- [14] P. G. Neumann, S. W. Moore, R. N. Watson, J. Anderson, N. Dave, B. Davis, J. H. Han, S. M. Hand, A. Horsman, M. Huxtable *et al.*, "Modular research-based composable trustworthy mission-oriented resilient clouds (mrc2)," SRI INTERNATIONAL Menlo Park United States, Tech. Rep., 2016.
- [15] J.-P. Nikkarila and M. Ristola, "'runet 2020'-deploying traditional elements of combat power in cyberspace?" in *2017 International Conference on Military Communications and Information Systems (ICMCIS)*. IEEE, 2017, pp. 1–8.
- [16] P. Tucker, "Russia plans to cut off some internet access next week," *Defense One*, 2019.
- [17] J. Wakefield, "Russia 'successfully tests' its unplugged internet," *BBC News*, 2019.
- [18] K. Bhasin and J. L. Hayden, "Space internet architectures and technologies for nasa enterprises," in *2001 IEEE Aerospace Conference Proceedings (Cat. No. 01TH8542)*, vol. 2. IEEE, 2001, pp. 2–931.
- [19] Y. Xu, J. Zhang, W. Wang, A. Juneja, and S. Bhattacharya, "Energy router: Architectures and functionalities toward energy internet," in *2011 IEEE International Conference on Smart Grid Communications (SmartGridComm)*. IEEE, 2011, pp. 31–36.
- [20] Q. Sun, R. Han, H. Zhang, J. Zhou, and J. M. Guerrero, "A multiagent-based consensus algorithm for distributed coordinated control of distributed generators in the energy internet," *IEEE transactions on smart grid*, vol. 6, no. 6, pp. 3006–3019, 2015.
- [21] N. Leveson, "Are you sure your software will not kill anyone?" *Communications of the ACM*, vol. 63, no. 2, pp. 25–28, 2020.
- [22] K. Zetter, *Countdown to Zero Day: Stuxnet and the launch of the world's first digital weapon*. Broadway books, 2014.
- [23] M. J. Kochenderfer, J. E. Holland, and J. P. Chryssanthacopoulos, "Next-generation airborne collision avoidance system," Massachusetts Institute of Technology-Lincoln Laboratory Lexington United States, Tech. Rep., 2012.
- [24] N. G. Leveson, *Safeware: system safety and computers*. ACM, 1995.
- [25] T. Ishimatsu, N. G. Leveson, J. Thomas, M. Katahira, Y. Miyamoto, and H. Nakao, "Modeling and hazard analysis using stpa," 2010.
- [26] I. Friedberg, K. McLaughlin, P. Smith, D. Laverty, and S. Sezer, "Stpa-safesec: Safety and security analysis for cyber-physical systems," *Journal of Information Security and Applications*, vol. 34, pp. 183–196, 2017.