# Cyber Situational Awareness of Critical Infrastructure Security Threats

Fatemeh Movafagh
*School of Computing Science*
*Simon Fraser University*
British Columbia, Canada
email: fma44@sfu.ca

Uwe Glässer
*School of Computing Science*
*Simon Fraser University*
British Columbia, Canada
email: glaesser@sfu.ca

***Abstract*—The rising frequency and sophistication of cyber-attacks pose a notorious threat to critical infrastructures, heavily reliant on industrial control systems for advanced automation. To explore this evolving challenge systematically, a robust cyber situational awareness framework is essential. Our paper adopts a dual approach, focusing on both the broader scope of threat mitigation and remediation to understand the breadth of the problem and on online intrusion detection applied to supervisory control data to comprehend its depth. The methodical framework and analytic model we propose here are tailored to cyber-physical systems used for industrial control and operational technology. By acknowledging transitional vulnerabilities in these systems, we stress the necessity of proactive measures to mitigate the risk of widespread cascading and escalating infrastructure failures. At the core of our contribution lies GenericAttackTracker, a novel analytic framework for online anomaly detection, which combines dynamic attack scoring with Bayesian inference to fuse results from supervisory control data analysis with real-time contextual information into actionable threat intelligence. By leveraging the abstract semantic properties of Heterogeneous Information Network Analysis for structural analysis and of Abstract State Machines for deriving executable abstract models of complex distributed systems, our framework supports a system of systems view of critical infrastructures and facilitates the daunting task of dynamically analyzing their intricate interdependencies.**

***Keywords*—*Cyber-physical systems; supervisory control systems; online threat detection; infrastructure interdependencies; machine learning; anomaly detection; dynamic attack scoring.***

## I. Introduction

Increasingly frequent and sophisticated cyberattacks have become a severe threat. Responding to the evolving cyber threat landscape, security technology is advancing, but not fast enough to keep pace with the threat. Security breaches frequently compromise the protection of sensitive information, exposing personal identities, intellectual property and financial assets. This trend means mounting damages in the hundreds of billions of dollars, erosion of trust in conducting business and collaboration in cyberspace and mounting fears of catastrophic events triggered by attacks that can physically cripple Critical Infrastructure (CI). Such attacks aim at indefinite disruptions of services that are essential for the functioning of our society and economy. In times of escalating political tensions and rising financial rewards from cybercrime, CI is at high risk from global cyber threat activity [1]–[3].

Critical infrastructures rely on Industrial Control Systems (ICS) as principal components of Operational Technology (OT) used for advanced automation of industrial processes. This includes different types of devices, systems, and networks to monitor and control physical processes, machinery, and other infrastructure components. Two standard control system architectures widely used for CI facilities are Supervisory Control and Data Acquisition (SCADA) and Distributed Control Systems (DCS) [4]. ICS technology offers robust and reliable solutions for advanced automation used in a variety of industries including manufacturing, oil and gas, electric energy generation and distribution, aviation, maritime, rail, and utilities, among many other CI sectors [5].

With progressive automation of critical industrial processes, the attack surface for sophisticated cyber threats expands, intensifying the risk of cascading and escalating failures [1][6]. When directly or indirectly connected to the Internet, ICS hardware and software can get exposed to illicit online access in attempts to exploit OT system vulnerabilities through various adversarial scenarios. A well-orchestrated cyberattack on a facility's integrated process control system may cause lasting and widespread disruptions and extensive physical damage by overloading vital system components [7]–[10]. Despite the many diverse uses, ICS architectures frequently build on the same core technologies, mostly SCADA, DCS, and Programmable Logic Controllers (PLC), for lower-level control tasks. SCADA systems and DCS are often networked together. Homogenous core architectures and tight coupling make them more vulnerable to cyberattacks because a single discovered vulnerability can potentially be exploited across several different systems [4].

This paper explores emerging threats to OT used in various critical sectors [5] and analyzes why CI security is a matter of growing concern that calls for enhanced resilience against the most aggressive threats. When vital components, systems, or networks get compromised, the incapacitation or destruction of CI assets could result in catastrophic loss of life, adverse economic effects and significant harm to public confidence [11]. The research presented here aims at a holistic methodical framework for devising a novel generic analytic model for cyber situational awareness of critical infrastructure threats. A central focus is on distributed online anomaly detection and interpretation of abnormal activity patterns in supervisory control data streamed from the operation of CI; i.e., patterns that deviate from the expected normal behavior beyond what could be explained by the presence of regular noise in the control data. The scope of our analytic model is not limited to single infrastructure entities but rather takes into account that multiple infrastructures are often interconnected as a system of systems with complex interdependencies [4]. "What happens to one infrastructure can, directly and indirectly, affect other infrastructures, impact large geographic regions, and send ripples throughout the national and global economy." [12]

The broader intent of our methodical framework is to also serve as a "lens" for gauging CI security and resilience against evermore advanced adversarial scenarios. Considering that network technology may never be completely secure, cybersecurity is about risk mitigation at the end of the day [13]. A holistic understanding of cybersecurity risks is crucial for making informed decisions on rational grounds. Risk mitigation strategies call for a complete assessment of vulnerabilities and consequential security risks to effectively enhance CI resilience. This fact was also stressed by the U.S. Government Accountability Office in their 2022 report on the U.S. electric grid security status: "DOE has developed plans to implement a national cybersecurity strategy for protecting the grid. However, we found that DOE's plans do not fully incorporate the key characteristics of an effective national strategy. For example, the strategy does not include a complete assessment of all the cybersecurity risks to the grid. Addressing this vulnerability is so important that we made it a priority recommendation for DOE to address." [14]–[16].

Besides the methodical framework, our main contribution is GenericAttackTracker, a distributed analytic framework for online detection and interpretation of anomalous activity patterns in supervisory control data. Building in part on our previous work, AttackTracker [17], the novel features of GenericAttackTracker significantly advance the core analytic model and expand the scope of AttackTracker to: (1) integrate contextual real-time threat intelligence and apply Bayesian inference to offer a broader decision basis and more reliable decision-making process; and (2) directly support a "system of systems" view for situational awareness of the cybersecurity status of multiple interdependent CI entities [18].

The remainder of this paper is organized as follows. Section II describes the broader scope of the problem in light of a multidimensional problem space. Next, Section III provides some background on industrial automation and the notorious challenges of online analysis and interpretation of supervisory process control data. Section IV introduces our methodical framework and generic analytic model. Building on AttackTracker, the generic model, GenericAttackTracker, expands the scope of the basic model in two principal ways. Finally, Section V concludes the paper.

## II. Problem Scope

The evolving threat landscape underscores the critical necessity for a robust cyber situational awareness framework. Such a framework should provide a comprehensive overview of a system's cyber environment, enabling quicker identification, understanding, and assessment of potential or existing threats, and mitigation approaches for such threats. This is particularly pertinent to OT and ICS, which oversee the functioning of CIs. Given their pivotal role in operating facilities such as electrical utilities, oil and gas pipelines, water utilities, chemical plants, and rail systems, among others [19], any compromise of these systems may result in serious disruptions and severe damage.

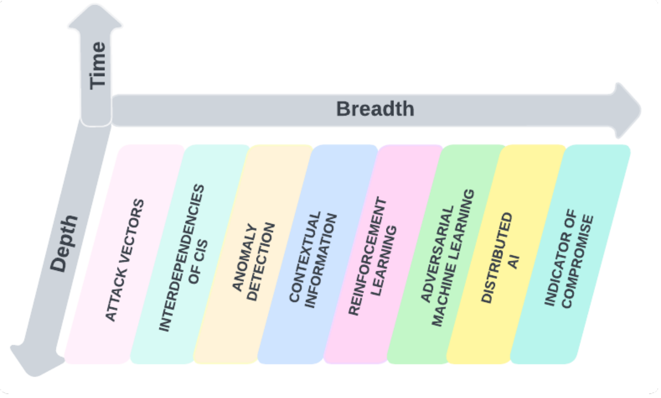

Figure 1. Multidimensional problem of security threats.

### A. Cyber Situational Awareness Framework

The multidimensional problem of security threats calls for a multifaceted solution, where multiple layers of defence mechanisms and strategies must be put in place to safeguard systems. Hence, to effectively tackle this intricate challenge, we propose a cyber situational awareness framework characterized by three dimensions: breadth, depth and time. This framework offers a holistic view of essential approaches for protecting ICS from escalating cyber threats, illustrated in Figure 1. As much as one must consider various aspects of defense breadth, one must also consider defense depth at the same time and routinely reassess both breadth and depth [1]. In fact, it is inadequate to defend in one dimension only. Defense that lacks depth despite breadth leaves vulnerabilities, while depth without breadth still allows attackers to find alternate entry points. Routine reassessment is critical to ensure that defense mechanisms remain fully intact and newly discovered vulnerabilities and exposures get patched in a timely manner. In the remainder of this section, we explore briefly some aspects of the breadth that also need to be taken care of in other dimensions.

*1) Attack Vector, Indicator of Compromise (IOCs):* OT is critical for industrial processes but exposes systems to cyber-attacks through various attack vectors, including network-based and physical process attacks [20]. Attackers employ advanced multi-vector strategies, targeting multiple entry points to exploit system vulnerabilities [21], emphasizing the need for a comprehensive and multifaceted defense approach to protect ICS. On the other hand, IOCs are forensic data logs that offer evidence of malicious activity on a system or network. Monitoring IOCs enables incident responders to detect signs of malicious actions and respond promptly to similar intrusions in their early stages [22].

*2) Interdependencies of CIs:* Due to the complex interdependencies between different infrastructures, a disturbance in one system can trigger cascading failures, leading to far-reaching and severe impacts. Thus, it becomes essential to leverage the data from one CI as a potential alert trigger for others. By doing so, we can anticipate and address the risk of cascading failures, strengthening the overall resilience of our critical systems. This part will be explored in depth in Sections III and IV.

*3) Anomaly Detection, Machine Learning, Reinforcement Learning, Contextual Information:* Anomaly detection particularly focused on time series data prevalent in ICS [23], forms a crucial aspect of the cyber situational awareness framework. Identifying temporal deviations in data patterns can be the key to uncovering potential threats or system malfunctions. Behavior-based and process-based anomaly detection are two approaches to safeguard ICS. The former uses machine learning to monitor system behavior and detect deviations, while the latter focuses on monitoring physical processes controlled by the ICS. These methods, augmented by reinforcement learning, address the limitations of traditional signature-based detection against novel and complex cyber-attacks such as zero-day attacks [24]. In addition, by integrating Bayesian inference, which utilizes probabilistic models, the detection process can dynamically update the likelihood of ongoing attacks based on incoming contextual data, and detect anomalous attack-based events accurately [25].

*4) Adversarial Machine Learning (AML):* The use of Machine Learning (ML), DL (Deep Learning), and RL (Reinforcement Learning) techniques in cybersecurity has improved threat detection. However, it also introduces vulnerabilities through AML attacks. These attacks can manipulate input data or the model itself to cause false positives and false negatives in anomaly detectors, weakening security system performance by exposing it to evasion and poisoning attacks [26]. The goal is to make ML, DL and RL in security a strength, and to enhance the resilience of OT and ICS security, not to be an exploitable vulnerability. Therefore, working on the robustness of anomaly detectors against AML such as what has been done in [27] is a must. This also demonstrates that staying ahead of threats requires constant situational awareness and readiness to respond to emerging cyber threats.

## III. Industrial Process Control

Automation is essential for the steady operation of critical infrastructure to continuously monitor and control machinery, systems, and processes; it enhances efficiency, productivity, quality of service delivery and safe operation of critical assets. We have thus become inexorably dependent on automated services and will be even more so with future smart industrial process control applications. An apt example to epitomize this ongoing trend is smart manufacturing in the fourth industrial revolution, tagged Industry 4.0 [28]. Under the cyber-physical system (CPS) paradigm, this situation is further exacerbated through the increasing integration of embedded computing with sensor networks (and other IoT devices) to monitor and control processes in the physical environment.

*1) Cascading and Escalating Failures:* While achieving great efficiencies through seamless interoperability of software and networking components with dynamics of physical processes, CPS technology intensifies fragility. When exploited by advanced threats, fragility amplifies the risk of cascading and escalating failures. Cascading events occur when local equipment failure or other disruptions trigger subsequent failures or disruptions on a larger scale.

Although triggered by "natural" causes, the phenomenon occurred in August 2003 for the Eastern Interconnection, one of the three major electric power grids in North America. A local fault of a high-voltage transmission line went unnoticed due to an alarm system malfunction, which in turn tripped a cascade of failures throughout southeastern Canada and eight northeastern U.S. states. In total, 50 million people lost power for up to two days in the biggest blackout in North American history [29]. In February 2021, the U.S. state of Texas suffered a major power crisis after severe winter storms, resulting in at least 246 deaths and property damages in excess of $\$195$B. Cascading failures propagated across multiple interdependent infrastructures causing insufficient power generation capacity online, which resulted in insufficient natural gas supply to the power plants. When power was cut, it disabled compressors that push gas through pipelines, knocking out further gas plants due to lack of supply [8].

*2) Abnormal Activity Patterns:* Critical processes require constant supervisory control of their operational status to issue alerts and initiate an emergency shutdown when abnormal activity patterns approach defined safety margins. Supervisory control data is temporal data interpreted as streamed sequences of real-value measurements taken at regular time intervals, referred to as time series data. Any observed activity patterns that do not conform to the expected behavior but seem to occur "out of place" are denoted as anomalies or outliers. An intuitive definition of the meaning of outlier is offered by Hawkins [30]: "an observation that deviates so much from other observations as to arouse suspicion that it was generated by a different mechanism."

*3) Anomaly Detection Challenges:* Real-world processes are notoriously prone to uncertainties caused by "external" factors such as communication errors, fluctuations in demand and supply, and technical instabilities resulting in inevitable variance in the data, characterized as noise. A number of factors make online anomaly detection in time series data streamed from the operation of a supervisory control system a challenging problem:

- Identifying anomalous activity patterns that often remain hidden to the human eye requires learning normal activity to train a robust model that not only fits previously observed data but also carries over to unobserved data; naturally, developing such a model is not a trivial task.
- Anomalies occur for various reasons, thus an even more intricate problem often is to differentiate the typically few anomalies of interest—above all, suspicious anomalous behavior indicating a potential security threat—from the vast majority of anomalies caused by noise, seasonality or other trends irrelevant to security.

Figure 2 illustrates common variance due to noise observed in time series data for electricity power consumption recorded at one datapoint per minute over four consecutive days. While

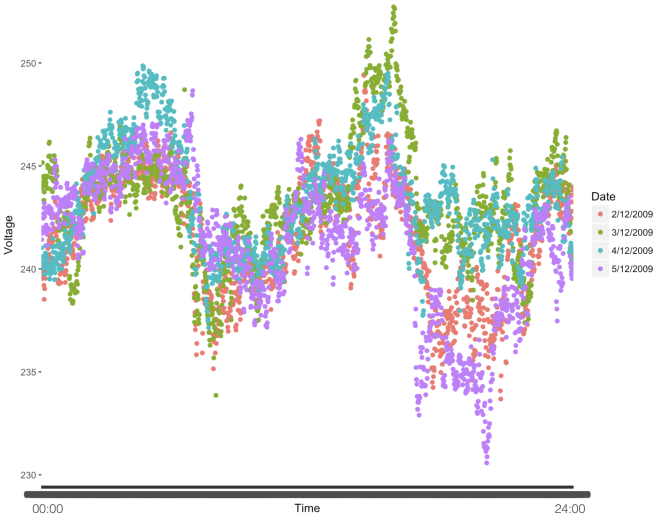

Figure 2. Electricity power consumption data for households over four 24-hour periods on consecutive weekdays show changes in voltage due to fluctuating demand and supply.

the exact power consumption behavior over a 24-hour time period differs on any given day, a recognizable overall pattern emerges; however, the boundaries for what constitutes normal variations of routine activity remain blurry. This phenomenon is persistent and not just due to the small sample size.

## IV. Methodological Framework

Online analysis of supervisory control data streamed from the operation of mission-critical systems is the basis for early threat activity detection. Finding suspicious and potentially harmful behavior anomalies without delay is key for swift mitigation and remediation to reduce the impact of an attack by launching countermeasures containing security breaches locally before they spread laterally across wider networks.

We first discuss AttackTracker, a distributed analytic model using dynamic attack scoring for online cyber threat activity detection for single infrastructure entities [17]. Building on this model, we then propose GenericAttackTracker, expanding the scope of the basic model to: (1) integrate contextual real-time threat intelligence and apply Bayesian inference for a broader and more reliable decision-making basis; and (2) support a "system of systems" view for situational awareness of multiple interdependent critical infrastructures [18].

### A. Attack Tracker

AttackTracker offers a robust and scalable framework based on a distributed analytical model for online tracing of threat activity patterns in supervisory control data by orchestrating a hierarchical network of threat activity detectors. This way, evidence of threat activity observed anywhere in the system is aggregated across control system architecture levels. Utilizing dynamic attack scoring boosts the analytic performance and reduces the false alarm rate by ignoring potential contextual noise and errors in the behavior prediction phase [17].

AttackTracker produces highly encouraging results [31] when applied to the Secure Water Treatment (SWaT) testbed created by Singapore University of Technology [32]. SWaT serves as a control signal source for data collected from a scaled-down version of an industrial water purification plant targeted by a variety of realistic attacks on different parts of the system. Figure 3 illustrates the basic AttackTracker architecture and its hierarchical organization.

A hierarchy of linked attack detectors continuously monitors the operation of controllers at different levels of a supervisory control system. $(l_1)$-detectors monitor peripheral controllers such as PLC units at the local level. At higher levels, $(l_i)$, for $i = 2, 3, \ldots$, detectors monitor the output of multiple detectors at level $(l_{i-1})$. At the top level, a single detector determines the global operational status of the whole system and reports attacks in progress in any one of the subsystems.

In other words, local detectors analyze the operation of a local subsystem as mirrored by the state of its sensors and actuators to spot abnormal patterns in the control data stream indicating a collective anomaly associated with an attack on this subsystem. Each local detector uses a Behavior Predictor module feeding the 'expected' next observation values into an Inference Engine. The Inference Engine module processes and labels observations, assigns attack scores, and raises red flags based on the deviation of observed values from predicted ones relative to a dynamically adjusted threshold. For $i \geq 2$, $(l_i)$-detectors aggregate data and information received from their lower-level detectors to determine the attack scope in the underlying levels. This way, detectors operating at higher levels are able to distinguish distributed threat activities in addition to centralized attacks.

For illustration, we consider a simple example in a SCADA-based testbed with three subsystems $(A, B, C)$. Subsystem $A$ has detected an unusual water pressure spike, while subsystem $C$ observed a decline in the water flow rate; no anomalies were found in subsystem $B$. Two secondary $l_2$-detectors overseeing pairs of these subsystems recognized the irregularities in $A$ and $C$. $A$ top-level detector, consolidating findings from the $l_2$-detectors, identified anomalies in two out of the three subsystems and signaled a system-wide alert. The top-level detector's Inference Engine deduced that the concurrent abnormalities in $A$ and $C$ might indicate a coordinated attack, recommending prompt action. This layered detection system ensures complex anomalies do not go unnoticed even when overlooked locally.

*1) Behavior Predictor:* Behavior Predictor is a core component that learns the normal behavior of a subsystem and forecasts the next local feature values based on previous observations. It uses a Multivariate Temporal Convolutional Network (MTCN) model to learn hidden patterns from a history of discrete observations in the form of a multivariate stochastic time series. Behavior Predictor is implemented to detect potential drifts in the stream data and adapt itself to not to be fooled by attacks.

*2) Inference Engine:* Inference Engine is the other component that decides based on a multi-modal view provided by its associated Behavior Predictor and the underlying detectors. It

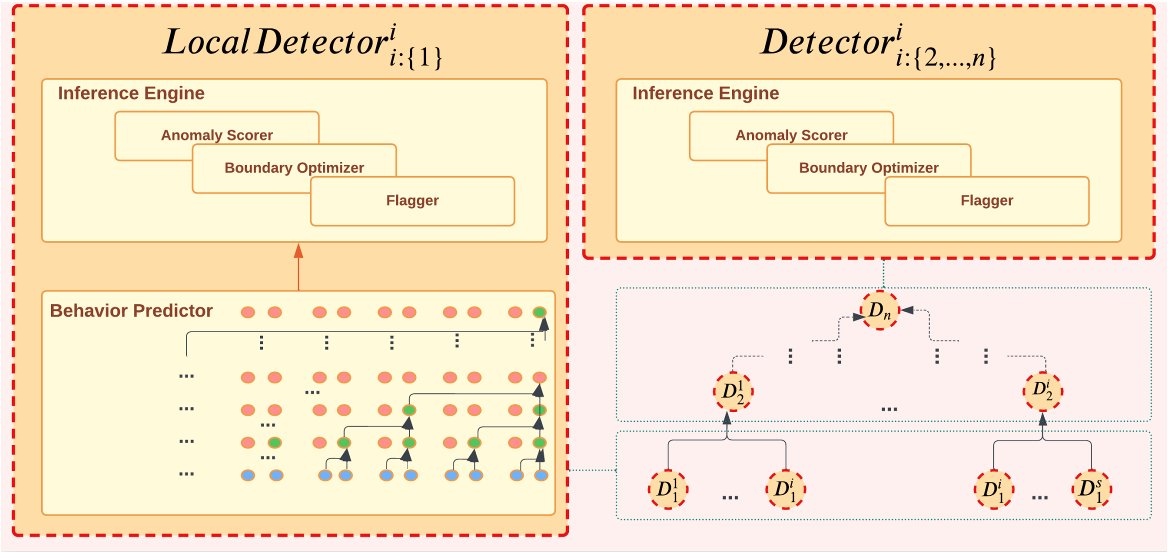

Figure 3. AttackTracker Framework Architecture: Local detectors operate at Level 1; regular detectors operate at all levels higher than Level 1.

enhances higher-level detectors by utilizing Behavior Predictor to trace the collective behavior of their underlying subsystems. Inference Engine is responsible for making decisions based on the information provided by the Behavior Predictor components and detectors, and it can help the end-user to choose the best mitigative action by highlighting the attack target and its potential cascading influences.

The Inference Engine component of the AttackTracker framework utilizes individual scoring and system-wide scoring as part of its anomaly detection process. The individual scoring phase involves offline and online steps, where a model is trained to detect anomalies based on transformed feature vectors. This phase focuses on detecting local subsystem anomalies. The system-wide scoring phase aggregates results from individual detectors to identify system-wide attacks, which individual detectors may miss due to the distributed nature of the network. Simultaneous anomalies and global log-based anomaly scores are considered in this phase.

The combination of individual and system-wide scoring enhances attack detection accuracy and reduces false alarms. This combination happens through the ”moving average” strategy in the Inference Engine component. It helps to identify collective and correlated anomalies as one single attack. This decision-making strategy is based on the trade-off between deviation and persistence, where a persistent anomalous interval is more suspicious of being an attack than a single strike caused by sensor noise or predictor faults.

*3) Dynamic Scoring:* The dynamic scoring method is based on a sliding window approach that considers the current observation and the previous observations to calculate the anomaly score. The anomaly score is then compared to a threshold value to determine whether an attack has occurred. The threshold value is dynamically adjusted based on the current state of the system and the historical data. The dynamic scoring method is designed to ignore potential contextual noise and errors in the Behavior Predictor components and to handle regular spikes of observed ”anomalies” in cases where they have not captured all the patterns of normal data.

### *B. Generic Attack Tracker*

Our GenericAttackTracker framework advances the analytic model and expands the scope of AttackTracker significantly by encompassing two principal novel features called: Bayesian View and System of Systems View (see below). Please note that this paper emphasizes the principles of GenericAttackTracker and their application, not the detailed implementation.

*1) Bayesian View:* A problematic aspect of time series anomaly detection in control data streamed from a mission-critical system is the rate of false positives: even when the relative rate is low, the absolute number of false positives may still be intolerable for high data volumes depending on a system’s critical mission. One way to mitigate the problem is fusing the results from control data analysis with contextual information from other potential sources of actionable threat intelligence to be used in the decision-making process. This leads to Bayesian methods. The strength of the Bayesian approach is its ability to combine information from multiple sources, thereby allowing greater ‘objectivity’ in final conclusions [33]. The result is a broader foundation for making more reliable decisions [34], whereas ignoring actionable threat intelligence originating from supplementary sources may come at the expense of missing out on the bigger picture.

A holistic view of threat activities calls for integrating data- with knowledge-driven threat analysis as a basis for applying Bayesian inference [35]. Assuming an evidential interpretation of probability, Bayes’ rule is used to update the probability for a hypothesis $H$ as more evidence or information $E$ becomes available. Formally, this is stated as a conditional probability:

$$P(H|E) = \frac{P(E|H) \cdot P(H)}{P(E)}, \text{ for } P(E) > 0, \tag{1}$$

where $P(E)$ is calculated as follows:

$$P(E) = P(H) \cdot P(E|H) + P(\neg H) \cdot P(E|\neg H) \tag{2}$$

In our case, $P(H|E)$ describes the probability of observing actual malicious activity as associated with a cyberattack based on prior knowledge of conditions that may be related to the event (abnormal patterns in the control data stream) before and after accounting for corroborative evidence from another threat intelligence source. A basic example is the Indicator Of Attack (IOA) status [36]: any digital or physical evidence that a cyberattack is likely imminent or in progress. IOAs generally focus on the intent of what an attacker is trying to accomplish, regardless of the malware or exploit used in an attack [37]. Events indicative of suspicious activity include: HTTP/HTTPS connections via non-standard ports (rather than port 80 or port 443); unusual network traffic; multiple user logins from different regions; internal hosts communicating with countries outside of the business range, among many others.

The following example illustrates the idea using numbers are not based on real-world or experimental data but are solely meant for the sake of explanation.

- Prior belief, $P(H)$: The value is determined by the Inference Engine component. Based on control data patterns, the anomaly detector estimates there's a 30% chance an attack is occurring: $P(H) = 0.30$.
- Evidence, $E$: The value is based on contextual information, for instance, an external threat intelligence feed that alerts us to an ongoing global cyber attack campaign.
- Likelihood, $P(E|H)$: This is the probability of receiving an external threat alert given that an attack is in progress. From past data, a value of $P(E|H) = 0.70$ is assumed.
- Probability of evidence, $P(E|\neg H)$: The probability of receiving an alert even when there is no attack is assumed to be $P(E|\neg H) = 0.10$ based on past data.

First, the probability of getting the external alert is calculated using Equation 2:

$$P(E) = 0.30 \times 0.70 + 0.70 \times 0.10 = 0.28$$

Next, the posterior (updated) probability of being under attack given the alert is calculated using Equation 1:

$$P(H|E) = \frac{0.70 \times 0.30}{0.28} \approx 0.75$$

Given the alert from the external threat intelligence feed, our updated belief that we're under attack went up from 30% (based solely on anomaly detection) to 75% (after accounting for the external threat intelligence).

Feeding supplementary IOA status information or threat intelligence from other alternative sources into the inference component of detector modules of GenericAttackTracker requires only a limited modification of AttackTracker's basic Inference Engine component. An example is threat intelligence derived from interdependencies between separate CIs. This reveals how an incident in one CI can ripple through and affect other CIs. A deeper exploration of this concept is provided in the following section.

*2) System of Systems View:* Generally, CI entities are highly interdependent in complex ways; an incident in one infrastructure can directly or indirectly affect related infrastructures, resulting in cascading and escalating failures [4]. The nature and reverberations of interdependencies are a complex and difficult problem to analyze. In their work, Rinaldi et al. [12] describe six dimensions of infrastructure interdependencies: types of interdependencies, infrastructure environment, coupling and response behavior, infrastructure characteristics, types of failures and state of operations. Although each has distinct characteristics, these classes of interdependencies are not mutually exclusive. Understanding these dimensions and applying them to the analysis of interdependencies among different CIs is crucial for maintaining a resilient system of systems. Incidents like the ransomware attack on the Colonial Pipeline in 2021 [38], or the large-scale electric grid failures cited in Section I, are vivid reminders that the impact due to interdependencies is very real.

Many studies of individual CI systems overlook their interconnection and mutual dependency [12]; only a few take interdependencies into account. However, these works either use simplified simulation platforms to analyze interdependencies among a limited type of CI entities, e.g., in [18], or they only measure risk based on interdependencies [39][40]. In contrast, GenericAttackTracker is designed to facilitate the modeling of CI systems with complex relations between their constituent entities. Our model abstractly identifies linked infrastructure entities as populations of interacting agents, in accordance with [12]. Nowadays, complex technical systems frequently comprise a large number of interacting, multi-typed components interconnected through communication and control networks. The information infrastructure of many such systems can abstractly be viewed as Heterogeneous Information Networks (HINs) [41] and be analyzed through Heterogenous Information Network Analysis (HINA) [42]. The HINA paradigm has gained wide attention from researchers in data mining and information retrieval fields; especially, it is used to mine hidden patterns through mining link relations from networked data [43].

The concepts of HIN/HINA align with our situation where multiple linked CIs with different interdependencies with each other are using GenericAttackTracker for online anomaly detection. GenericAttackTracker enhances our methodological framework by utilizing HINA [44][43]. This way, we model and analyze static representations of CI interdependencies for a more realistic approach to anomaly detection. With this objective, we define the following sets as interdependency categorises [12]:

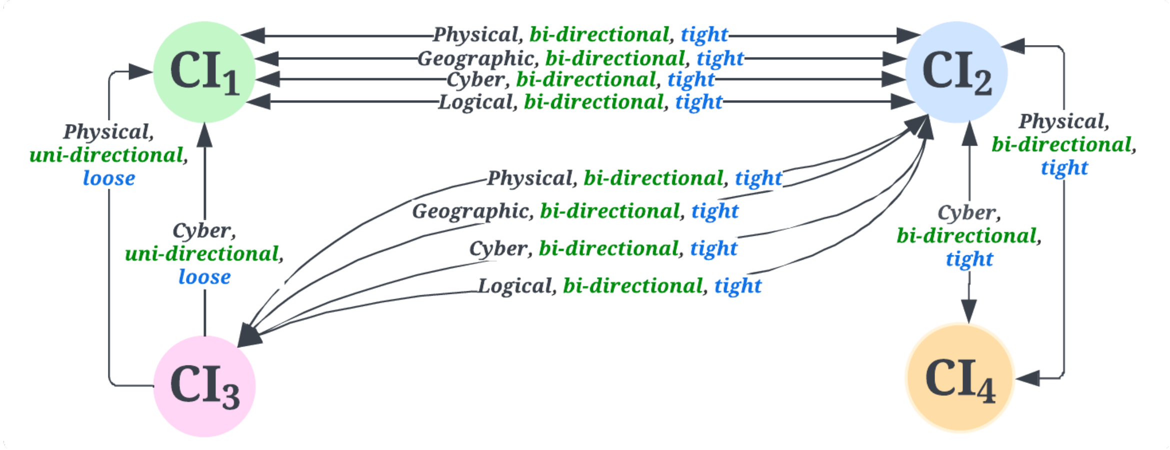
CI$_1$
CI$_2$
CI$_3$
CI$_4$
Physical, bi-directional, tight
Geographic, bi-directional, tight
Cyber, bi-directional, tight
Logical, bi-directional, tight
Physical, uni-directional, loose
Cyber, uni-directional, loose
Physical, bi-directional, tight
Geographic, bi-directional, tight
Cyber, bi-directional, tight
Logical, bi-directional, tight
Physical, bi-directional, tight
Cyber, bi-directional, tight

Figure 4. The Network schema $NS_G$ comprises four different infrastructure types; directional links between the various node types state constraints on CI interdependencies.

$Class = \{Physical, Cyber, Geographic, Logical\}$
$Direction = \{uni-directional, bi-directional\}$
$Degree = \{loose, tight\}$

We build on the HIN definition in [43]: A heterogeneous information network $G(V, E, A, R)$ is composed of an object set $V = \{n_1, n_2, ..., n_n\}$ with object types $A = \{a_1, a_2, ..., a_m\}$, and a set of links $E = \{e_1, e_2, ..., e_k\}$ with relation types $R = \{r_1, r_2, ..., r_l\}$, where $|A| > 1$ or $|R| > 1$ (to differentiate HINs from regular graphs). Two surjective function mappings assign object types to objects and relation types to links. If two links belong to the same relation type, the two links share the same starting object type as well as the ending object type.

For a better understanding of the composition of a complex HIN $G$, the network schema $NS_G$ is a template that describes the meta network structure of $G$ by specifying type constraints on the sets of objects and links of $G$. The result is a directed graph defined over object types $A$, with edges that are relation types from $R$. An HIN that conforms with a network schema is called a network instance of the schema. For a link type $R$ connecting object type $S$ to object type $T$, denoted by $S \xrightarrow{R} T$, S and T are the source and target object type of link type R.

A meta path $\mathcal{P}$ is a path defined on a schema $S_G = (\mathcal{A}, \mathcal{R})$, and is denoted in the form of $A_1 \xrightarrow{R_1} A_2 \xrightarrow{R_2} ... \xrightarrow{R_l} A_{l+1}$, which defines a composite relation $R = R_1 \circ R_2 \circ ... \circ R_l$ between objects $A_1, A_2, ..., A_l$, where $\circ$ denotes the composition operator on relations. The rich semantics of meta path is an important characteristic of HIN. Based on different meta paths, objects have different connection relations with diverse path semantics, which may affect many data mining tasks including clustering, classification, link prediction, ranking, and information fusion [43].

Figure 4 is an example of a HIN graph, $G$, that shows a network schema, $NS_G$, of four SCADA-based CI entities and their interdependencies derived from the NIST guide to ICS security [4]. These CIs are natural gas pipelines, electric power grids, water distribution systems, and railway transportation systems. Natural gas pipelines need electric power for their compressors, storage and control systems. On the other hand, electric power generation needs natural gas as a main or backup fuel for its generators. Thus, the physical interdependencies between these two CI types are bidirectional. Natural gas pipelines also might need water from case to case for cooling or emission reduction. So this is a unidirectional and loose physical interdependency.

Not only is electric power supply essential for the operation of railway transportation and water distribution systems but these two CI types are essential for power generation. Hence, the physical interdependencies between them are bidirectional. Within the GenericAttackTracker framework, each of these CI types acts as an agent that interacts with other CI entities. Bidirectional cyber interdependencies must be considered for all interdependent CIs. Building upon the graph in Figure 4, it is plausible that a disruption in the natural gas infrastructure can cause power disruptions, and electric power failures may lead to disruptions in other infrastructures.

In Figure 5, we delve deeper into a specific instance of the $NS_G$ of Figure 4, where the CIs are not just represented in their general form. Indeed, by identifying and analyzing different meta paths $\mathcal{P}$ within this schema graph, we can undertake a range of data mining tasks. Each unique meta path reveals distinct insights into the intricate interdependencies existing among the CIs.

Finally, it is not necessary to manually produce complex HIN graph structures as these can be generated automatically through the use of representation learning methods [45]. The details are beyond the scope of this paper.

Going beyond identifying and understanding normally static interdependencies, the final challenge is how to operationalize the cyber situational awareness framework and analytic model as needed for determining the broader impact of dynamically cascading and escalating failure scenarios in a timely manner. The state of operation of an infrastructure can be thought of

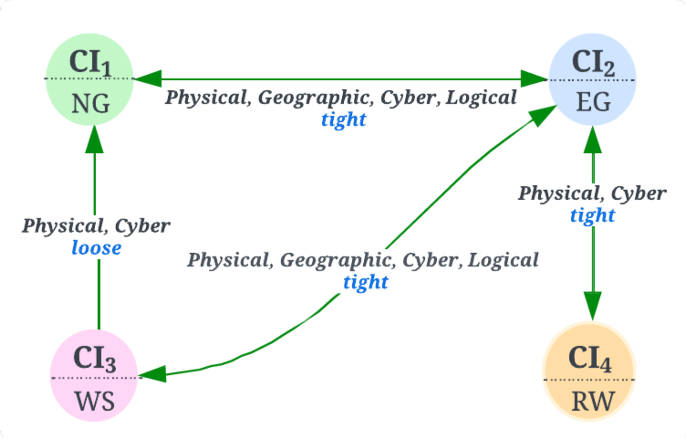
CI1
NG
CI2
EG
CI3
WS
CI4
RW
Physical, Geographic, Cyber, Logical
tight
Physical, Cyber
loose
Physical, Geographic, Cyber, Logical
tight
Physical, Cyber
tight

Figure 5. Network instance of schema $NS_G$: $i$) Natural gas pipeline (NG), $ii$) Electric grid (EG), $iii$) Water distribution systems (WS), $iv$) Railway transportation systems (RW).

as a continuum that exhibits different behaviors during normal operation conditions, times of severe stress or disruption, or repair and restoration activities. At any point in the continuum, the state of operation is a function of interrelated factors and system conditions [12]. This may be the hardest task after all.

Any viable solution does require continual reassessment of the security status of complex CIs and their constituent entities to account for emerging cyber threat events and incidents. By viewing linked infrastructure entities as interacting agents, the impact of threat activities on the operational status of related entities is modeled in terms of a distributed abstract machine. The model computes the situational awareness status of a complex CI based on the combined status of the component CI entities. The underlying formal model for developing the abstract machine builds on the operational modeling paradigm of Abstract State Machines (ASM) [46] and its method for stepwise refinement [47].

A distributed ASM, by definition, is a collection of asynchronously interacting ASM agents that collectively update a distributed global state. In previous work, we have successfully used the ASM paradigm as formal semantic foundation for modeling a complicated distributed situation analysis framework for maritime security [48] and also designed and developed a computational platform for making such models executable [49]. The design of the CI abstract machine model exceeds the scope of this paper but will be the subject of a separate paper.

Finally, in a system-of-systems context, the dynamic analysis of the operational status of interacting CI entities can also generate threat intelligence as input for the Inference Engine of GenericAttackTracker to enhance anomaly detection accuracy and overall system resilience. An attack on one entity may also spell trouble for other interdependent entities downstream.

## V. Conclusion and Future Work

With evermore sophisticated and damaging threats targeting critical infrastructure, cyber risks are intensifying and security breaches are more and more inevitable. In light of expanding the attack surface for advanced threats and zero-day exploits, enhancing the resilience of operational technology against the most serious threats is critical. Risk mitigation strategies call for a complete assessment of vulnerabilities and consequential risks to make informed decisions for effective risk mitigation and remediation on rational grounds.

Our main technical contribution, GenericAttackTracker, is a distributed and scalable analytic framework for detection of threat activity patterns in supervisory control data; its novel features significantly advance and expand the scope of the analytic core of the basic AttackTracker model in two principal ways: 1) fusing results from control data analysis with contextual threat intelligence from IOA sources into actionable insights yields a broader, more reliable decision basis, expected to further reduce false positive rates; 2) modeling infrastructures as interacting agents linked in complex ways—both physically and through ICT, i.e., what happens to one infrastructure can directly or indirectly affect other infrastructures—supports a “system of systems” view for situational awareness of the security status of multiple interdependent infrastructure entities.

Handling CI security is a complex and challenging task. In our research, we tackle this problem by combining advanced analytic methods with intuitive modeling paradigms to manage complexities. The ultimate goal is a coherent and consistent integration in an abstract methodical framework that facilitates a holistic view of the full scale problem scope.

While putting a spotlight on SCADA, a prevalent industry standard for monitoring and control of vital services not only in North America, the strategies we discuss here do likely apply to a much broader range of industrial process control systems. By exploring the feasibility of our approach for SCADA architectures, we aim to show the practical relevance of our analytic framework for ICS/OT at large.

Our future work, will continue the research to model and analyze complex network schemas of linked infrastructures to extract and interpret more intricate interdependencies. We believe HIN/HINA provides the expressiveness needed to tackle these tasks. Further, we will build upon our previous work on modeling distributed situation analysis processes as Abstract State Machine models, focussing on dynamically evaluating the status of complex CIs in near real-time.

## Acknowledgment

The research project for the work presented here was funded by the Natural Sciences and Engineering Research Council of Canada (NSERC) and by Innovation, Science and Economic Development Canada (ISED).

## References

[1] O. S. Saydjari, “Engineering trustworthy systems: a principled approach to cybersecurity,” *Communications of the ACM*, vol. 62, pp. 63–69, May 2019.

[2] CrowdStrike, “CrowdStrike 2023 Global Threat Report.” Available upon online request: https://go.crowdstrike.com/2023-global-threat-report. Accessed: 2023.08.31.

[3] Canadian Center for Cyber Security, “National Cyber Threat Assessment 2023–2024,” *Communications Security Establishment*, 2022.

[4] K. Stouffer, V. Pillitteri, S. Lightman, M. Abrams, and A. Hahn, “Guide to Industrial Control Systems (ICS) Security,” Tech. Rep. NIST SP 800-82r2, National Institute of Standards and Technology, June 2015.

[5] Cybersecurity & Infrastructure Security Agency, U.S. Department of Homeland Security, “Critical Infrastructure Sectors.” Online: https://www.cisa.gov/topics/critical-infrastructure-security-and-resilience/critical-infrastructure-sectors. Accessed: 2023.08.01.

[6] D. W. Hubbard and R. Seiersen, “How to Measure Anything in Cybersecurity Risk,” *John Wiley & Sons*, 2016.

[7] T. Tsvetanov and S. Slaria, “The effect of the colonial pipeline shutdown on gasoline prices,” *Economics Letters*, vol. 209, p. 110122, 2021.

[8] K. Madhavan and D. Rajamani, “2021 Texas Electricity Black-out Crisis: Root-cause Analysis and Recommendations,” *Journal of Student Research*, vol. 11, no. 1, pp. 1–10, 2022.

[9] A. Matrosov, E. Rodionov, D. Harley, and J. Malcho, “Stuxnet under the microscope,” *ESET LLC (September 2010)*, vol. 6, pp. 1–85, 2010.

[10] I. Thomson, “Everything you need to know about the petya, er, notpetya nasty trashing pcs worldwide.” Online: https://www.theregister.com/2017/06/28/petya_notpetya_ransomware/, 2017. Accessed: 2023.08.31.

[11] Public Safety Canada, “Canada’s Critical Infrastructure.” Online: https://www.publicsafety.gc.ca/cnt/ntnl-scrt/crtcl-nfrstrctr/cci-iec-en.aspx, 2022. Accessed: 2023.08.01.

[12] S. M. Rinaldi, J. P. Peerenboom, and T. K. Kelly, “Identifying, understanding, and analyzing critical infrastructure interdependencies,” *IEEE control systems magazine*, vol. 21, no. 6, pp. 11–25, 2001.

[13] T. Bossert and U.S. Department of Homeland Security, “Presentation at Cyber Week 2017,” *Blavatnik Interdisciplinary Cyber Research Center, Tel Aviv University, Israel*, 2017.

[14] U.S. Government Accountability Office, “Securing the U.S. Electricity Grid from Cyberattacks.” Online: https://www.gao.gov/blog/securing-u.s.-electricity-grid-cyberattacks, Oct. 2022. Accessed: 2023.08.31.

[15] U.S. Government Accountability Office, “ElectricityL Grid Cybersecurity: DOE Needs to Ensure Its Plans Fully Address Risks to Distribution Systems.” Online: https://www.gao.gov/assets/720/713257.pdf, Mar. 2021. Accessed: 2023.08.31.

[16] U.S. Government Accountability Office, “Critical Infrastructure Protection: Actions Needed to Address Significant Cybersecurity Risks Facing the Electric Grid.” Online: https://www.gao.gov/assets/710/701114.pdf, Aug 2019. Accessed: 2023.08.31.

[17] Z. Zohrevand and U. Glässer, “Dynamic attack scoring using distributed local detectors,” in *ICASSP 2020-IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 2892–2896, 2020.

[18] I. Eusgeld, C. Nan, and S. Dietz, “System-of-Systems Approach for Interdependent Critical Infrastructures,” *Reliability Engineering & System Safety*, vol. 96, no. 6, pp. 679–686, 2011.

[19] M. Conti, D. Donadel, and F. Turrin, “A Survey on Industrial Control System Testbeds and Datasets for Security Research,” *IEEE Communications Surveys & Tutorials*, vol. 23, no. 4, pp. 2248–2294, 2021.

[20] T. Mason and B. Zhou, “Digital forensics process of an attack vector in ics environment,” in *2021 IEEE International Conference on Big Data (Big Data)*, pp. 2532–2541, IEEE, 2021.

[21] E. Irmak and İ. Erkek, “An overview of cyber-attack vectors on scada systems,” in *2018 6th international symposium on digital forensic and security (ISDFS)*, pp. 1–5, IEEE, 2018.

[22] M. Asiri, N. Saxena, R. Gjomemo, and P. Burnap, “Understanding indicators of compromise against cyber-attacks in industrial control systems: a security perspective,” *ACM transactions on cyber-physical systems*, vol. 7, no. 2, pp. 1–33, 2023.

[23] B. Kim et al., “A comparative study of time series anomaly detection models for industrial control systems,” *Sensors*, vol. 23, no. 3, p. 1310, 2023.

[24] MR. Gauthama Raman, C. M. Ahmed, and A. Math, “Machine learning for intrusion detection in industrial control systems: challenges and lessons from experimental evaluation,” *Cybersecurity*, vol. 4, pp. 1–12, 2021.

[25] C. Wang et al., “Robust intrusion detection for industrial control systems using improved autoencoder and bayesian gaussian mixture model,” *Mathematics*, vol. 11, no. 9, p. 2048, 2023.

[26] S. Zhou et al., “Adversarial attacks and defenses in deep learning: From a perspective of cybersecurity,” *ACM Computing Surveys*, vol. 55, no. 8, pp. 1–39, 2022.

[27] Y. Jia et al., “Adversarial attacks and mitigation for anomaly detectors of cyber-physical systems,” *International Journal of Critical Infrastructure Protection*, vol. 34, p. 100452, 2021.

[28] L. Thames and D. Schaefer, “Industry 4.0: An overview of key benefits, technologies, and challenges,” *Cybersecurity for Industry 4.0: Analysis for Design and Manufacturing*, pp. 1–33, 2017.

[29] J. Minkel, “The 2003 northeast blackout–five years later,” *Scientific American*, vol. 13, pp. 1–3, 2008.

[30] D. M. Hawkins, *Identification of outliers*, vol. 11. Springer, 1980.

[31] Z. Zohrevand, *End-to-end anomaly detection in stream data*. PhD thesis, School of Computing Science, Simon Fraser University, 2021.

[32] J. Goh et al., “A dataset to support research in the design of secure water treatment systems,” in *Critical Information Infrastructures Security: 11th International Conference, CRITIS 2016, Paris, France, October 10–12, 2016, Revised Selected Papers 11*, pp. 88–99, Springer, 2017.

[33] A. Gelman et al., *Bayesian Data Analysis ($3^{rd}$ Edition)*. CRC Press, 2014.

[34] N. Silver, *The signal and the noise: the art and science of prediction*. Penguin UK, 2012.

[35] R. Kelter, “Statistical Rethinking: A Bayesian Course with Examples in R and STAN,” 2020.

[36] E. Kost, “What are IOAs? How they differ from IOCs.” Online: https://www.upguard.com/blog/what-are-indicators-of-attack, April 2023. Accessed: 2023.08.31.

[37] CrowdStrike, “IOA VS IOC.” Online: https://www.crowdstrike.com/cybersecurity-101/indicators-of-compromise/ioa-vs-ioc/, October 2022. Accessed: 2023.08.31.

[38] C. Wilkie, “Colonial Pipeline paid $5 million ransomware one day after cyberattack, CEO tells Senate.” Online: https://www.cnbc.com/2021/06/08/colonial-pipeline-ceo-testifies-on-first-hours-of-ransomware-attack.html, 2021. Accessed: 2023.08.01.

[39] A. O. Adetoye, M. Goldsmith, and S. Creese, “Analysis of dependencies in critical infrastructures,” in *Critical Information Infrastructure Security: 6th International Workshop, 2011, Lucerne, Switzerland, September 8-9, 2011, Revised Selected Papers 6*, pp. 18–29, Springer, 2013.

[40] P. Kotzanikolaou, M. Theoharidou, and D. Gritzalis, “Interdependencies between critical infrastructures: Analyzing the risk of cascading effects,” in *Critical Information Infrastructure Security: 6th International Workshop, CRITIS 2011, Lucerne, Switzerland, September 8-9, 2011, Revised Selected Papers 6*, pp. 104–115, Springer, 2013.

[41] J. Han, “Mining heterogeneous information networks by exploring the power of links,” in *International conference on discovery science*, pp. 13–30, Springer, 2009.

[42] Y. Sun and J. Han, “Mining heterogeneous information networks: a structural analysis approach,” *Acm Sigkdd Explorations Newsletter*, vol. 14, no. 2, pp. 20–28, 2013.

[43] C. Shi, Y. Li, J. Zhang, Y. Sun, and S. Y. Philip, “A survey of heterogeneous information network analysis,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 29, no. 1, pp. 17–37, 2016.

[44] C. Shi, SY. Philip, *Heterogeneous Information Network Analysis and Applications*. Springer, 2017.

[45] Y. Lei, L. Chen, Y. Li, R. Xiao, and Z. Liu, “Robust and fast representation learning for heterogeneous information networks,” *Frontiers in Physics*, vol. 11, p. 1196294, 2023.

[46] E. Börger and R. Stärk, *Abstract State Machines: A Method for High-Level System Design and Analysis*. Springer, 2003.

[47] E. Börger, “The ASM refinement method,” *Formal Aspects of Computing*, vol. 15, no. 2, pp. 237–257.

[48] N. Nalbandyan, U. Glässer, H. Y. Shahir, and H. Wehn, “Distributed situation analysis: A formal semantic framework,” in *Abstract State Machines, Alloy, B, TLA, VDM, and Z: 4th International Conference, ABZ 2014, Toulouse, France, June 2-6, 2014. Proceedings 4*, pp. 158–173, Springer, 2014.

[49] R. Farahbod, V. Gervasi, and U. Glässer, “Executable formal specifications of complex distributed systems with CoreASM,” *Science of Computer Programming*, vol. 79, pp. 23–38, 2014.