Analyzing Power Grid, ICT, and Market Without Domain Knowledge Using Distributed Artificial Intelligence

Eric MSP Veith[1], Stephan Balduin[1], Nils Wenninghoff[1], Martin Tröschel[1], Lars Fischer[1], Astrid Nieße[2], Thomas Wolgast[2], Richard Sethmann[3], Bastian Fraune[3], Torben Woltjen[3]

[1] OFFIS e.V.
R&D Division Energy
Oldenburg, Germany
Email:
first.last@offis.de

[2] Carl von Ossietzky University
Institute for Digitalized Energy Systems
Oldenburg, Germany
Email: first.last@uol.de

[3] Hochschule Bremen
Department for Computer Networks and Information Security
Bremen, Germany
Email:
first.last@hs-bremen.de

***Abstract*—Modern Cyber-Physical Systems (CPSs), such as our energy infrastructure, are becoming increasingly complex: An ever-higher share of Artificial Intelligence (AI)-based technologies use the Information and Communication Technology (ICT) facet of energy systems for operation optimization, cost efficiency, and to reach CO_2 goals worldwide. At the same time, markets with increased flexibility and ever shorter trade horizons enable the multi-stakeholder situation that is emerging in this setting. These systems still form critical infrastructures that need to perform with highest reliability. However, today's CPSs are becoming too complex to be analyzed in the traditional monolithic approach, where each domain, e.g., power grid and ICT, as well as the energy market, are considered as separate entities while ignoring dependencies and side-effects. To achieve an overall analysis, we introduce the concept for an application of distributed artificial intelligence as a self-adaptive analysis tool that is able to analyze the dependencies between domains in CPSs by attacking them. It eschews pre-configured domain knowledge, instead exploring the CPS domains for emergent risk situations and exploitable loopholes in codices, with a focus on rational market actors that exploit the system while still following the market rules.**

***Keywords*—Cyber-Physical Systems Analysis; Distributed Artificial Intelligence; Reinforcement Learning; ICT Security; Market Design.**

I. Introduction

During the last two decades, the power grid has seen an enormous development in the adoption of Information and Communication Technology (ICT) on a large scale in order to facilitate the inclusion of advanced methodologies, including Artificial Intelligence (AI)-based approaches. This increases efficiency and flexibility, which ultimately allows a higher share of renewable energy sources in the grid. However, together with a proceeding decentralization and the inclusion of energy markets, the complexity of the overall system also increased, with different factors adding to it, e.g., prosumers directly selling their Photovoltaic (PV) power or new market-based concepts for ancillary service provisioning, which need to be implemented by 2021 as per EU regulations [1].

Decentralized generation and consumption has led to the emergence of decentralized grid operation and control paradigms, many of which feature independent software agents. These Multi Agent Systems (MAS) exist for different tasks, e.g., to equalize real power generation and consumption, or to facilitate voltage control on local levels. A newer example of such a decentralized, specifically all-encompassing MAS that is aimed at including a high share of volatile, renewable energy sources is the *Universal Smart Grid Agent* system [2]–[4].

Assuming that major Internet of Things (IoT) trends will also influence the future power grid, the comprehensive use of ICT and AI technologies will, through their complexity, inevitably create an obstacle for a reliable operation of the power grid [5], [6]. At least since the cyber attack on the power grid of the Ukraine in December 2015 [7], [8], energy systems are recognized as valuable and vulnerable targets. Further attacks were seen in different stages with varying targets until and beyond 2017 [9]. These attacks demonstrate how ICT has a vital role in modern energy distribution networks. It needs to be reliable to ensure a stable power grid. However, due to the increasing ICT in modern power grids, the attack surface is getting bigger. Darknet marketplaces offer Distributed Denial-of-Service (DDoS)-as-a-Service and other attack-services for small money [10], which demonstrates that security testing is getting more important in this special domain.

Research actively addresses the numerous challenges that arise from the increased complexity and, thus, new attack vectors the emerge not only in the energy domain, but all Cyber-Physical Systems (CPSs) in general. Among them are neural control falsification, e.g., through Adversarial Learning (AL) [11]–[13], false data injection as attacks on state estimators [14]–[18], or utilizing compromised assets to actively damage the CPS [19].

In addition, a new type of attack has emerged in market-connected CPS like energy systems: The attack as a side effect of economically rational behavior. Energy markets are highly regulated in all countries. The need for regulation directly follows from the energy systems' inherent dependability on a dedicated infrastructure, like power grids, gas and heat networks. With this kind of infrastructure, a natural monopoly

 ISBN: 978-1-61208-818-1

is given. To ensure system stability while optimizing costs, market-based approaches are regulated to realize access to this infrastructure and system stability responsibility. The adaption of regulative frameworks is late by design: Once a loophole has been found, regulation is readjusted. Even if no outright cyber-attack is staged, actors in the market might exploit loopholes while still conforming to the rules. There are a couple of known examples where this has been done and actually affected the power grid, e.g., in Germany with Inc-Dec Gaming against the zonal system with uniform pricing scheme [20], or in another case in Great Britain [21].

However, in a recent survey looking at CPSs from the perspective of AI research, we found that a large portion of research focuses on a safe inclusion of AI technologies, such as Deep Learning or decentralized control through MAS in critical infrastructures, but also emphasizes the gaps between almost fully analyzed, reliable CPS and the complexity introduced by these techniques. Additionally, there is currently no systemic analysis approach that includes AI technologies as the driver to explore and analyze unknown CPSs for safety [22]. This survey can be seen as the main motivational background for this work: Traditional methods for analyzing the operational safety of a CPS can only cover specific, partial aspects. Hence, we found extensive research into many different aspects of safe CPS operation, but no approach for systemic testing of intra- and inter-domain relationships. From the point of the analysis, this causes a fragmentation of the whole system into islands. Aggregating subsystems also means that the effects of the interaction of components, as well as the influence of market actors is not completely covered. This holds especially true for systemic vulnerabilities, in which isolated parameters are within nominal boundaries, but the overall system is being destabilized through emergent effects. On the basis of the challenges outlined above, we create an intelligent, cross-sectional software technology for analyzing complex CPS in project PYRATE. It analyzes complex CPS with interdependent components autonomously, finding vulnerabilities leading to systemic failures. The core of the software technology to be developed is based on learning software agents that interact with a model—ideally a digital twin—of a CPS, using the resulting system states as reinforcing feedback signals for full self-adaptivity to efficiently explore the search space of actions for destabilizing ones.

Our project works on two different levels: On a methodical level, we plan to develop a universal methodology to analyze weaknesses of arbitrary CPS by finding successful attack strategies. On a practical level, we apply this methodology to an exemplary scenario containing a power system, an ancillary service market, and an ICT system, to demonstrate possible applications and the effectiveness of the methodology.

The remainder of this paper is structured as follows: Due to didactic reasons, in Section II, we first introduce the three environments of our demonstrator, explain major challenges in them, and describe the co-simulation setup. Afterwards, in Section III, we follow with a description of our cross-domain learning MAS that explores a CPS in order to defeat it. The experimentation process that underpins any analysis of our technology is described next, in Section IV, followed by the post-run analysis in Section V that aims to isolate the minimal chain of actions that led to CPS failure. Finally, in Section VI, we conclude with an outlook towards the realization.

II. Environment Under Scrutiny: A Demonstrator

In the research project, a power grid, an ICT network, and a local ancillary service market are simultaneously subjected to analysis, since the goal is to analyze interdependent behavior. Since the analysis cannot be performed on real infrastructure for obvious reasons, simulation models of each of the different domains are being synchronized at run-time using a co-simulation approach.

A. Power System

In this project's demonstrator, we focus on distribution grids to show the feasibility of the approach. Today's distribution grids lend themselves very well: They contain both, distributed large and aggregateable small loads, connect the major portion of Distributed Energy Resources (DERs), and are currently subject to large-scale ICT inclusion, as well as the development of local ancillary market concepts. Furthermore, they form the smallest meaningful, mostly self-contained environment that features a complex CPS with a variety of outside influence factors such as volatile power generation from renewable energy sources.

For simulation and benchmark purposes on distribution grid level, a scenario-based benchmark environment was developed. This benchmark environment incorporates a Medium Voltage (MV) grid developed by the International Council on Large Electric Systems (CIGRE) [23], [24], time series data of one year in 15min resolution (e.g., for wind, solar radiation, or consumption) from a former research project *Smart Nord* [25], and different component models, like PV or Combined Heat and Power (CHP).

B. Ancillary Service Market

For current energy markets, regulation is mainly settled, though adaptions still can be seen quite often, e.g., for optimization reasons. When implementing new energy markets, a whole new set of regulations is needed, though: There is a lot of activity in the implementation of regional energy markets and cell-based approaches, which are still in their infancy. Thus, we can expect many upcoming iterations on the regulation sets [26]. This holds especially true for all kinds of ancillary service markets, e.g., reactive power or flexibility markets [27], [28].

In this context, even new problems arise: We found that, for grid-stabilizing ancillary service markets, regional actors and even private households could cooperatively induce problems into the grid to later get paid for eliminating these very problems. E.g., if we assume that the grid operator has to procure reactive power in a purely market-based way, private households could synchronize their load behavior in order to manipulate the local voltage level and to violate the voltage band. That forces the grid operator to announce a reactive power

 ISBN: 978-1-61208-818-1

auction, in which generator agents would offer reactive power provision as ancillary service. Afterwards, the generator and household agents would divide profits and start a new attack. Regulation for such problems is not known at all, especially as this kind of malicious behavior is difficult to detect and proof.

Our methodology will help to systematically investigate and understand such profit-driven attacks, which will in turn allow for better market designs. For this, a local auction-based reactive power market with simple rules will be implemented as incentive for profit-driven attacks. This will allow for better understanding of possible attack vectors for profit maximization. Later, systematic comparison with more sophisticated market designs and rules will enable insights which market rules increase resilience against which attack strategies. Finally, we hope to find market designs that minimize attacks and that maximize grid stability, as well as detectability of such attacks.

C. ICT Simulation

Distributed power units are equipped with ICT to connect to wide-area networks. This enables operators to regulate and monitor distributed locations remotely, which is the foundation for implementing local ancillary service markets at the distribution grid level. The CIGRE MV model only specifies a distribution grid topology without covering the ICT domain, thus, we extend and overlay it with relevant ICT components to model a realistic multi-domain distribution grid infrastructure. Consequently, each node of the energy grid is accompanied by the corresponding representation in the ICT network that would, in reality, provide access to relevant sensors and actuators. Additionally, a communication network is built with routers and switches that connects these CPSs, arranged in multiple subnets, hence modelling a realistic ICT network.

Specific requirements arise from the multi-domain co-simulation setting: First, it needs to be efficient at simulating large networks. Second, the ICT simulation is required to create an accurate model of the reality and, therefore, compute realistic results, which is especially important when examining networks in a security context. Thus, it is necessary that existing software can be integrated with minimal modifications. Lastly, the simulation tool needs to be easy-to-use, so that also experts of the other simulated domains—who might have limited knowledge about ICT networks—can work with it after a short period of time. As there is no such simulator available that can meet all of these requirements, the *rettij* network simulator was developed. It is designed to simulate ICT components like routers, switches, clients and servers, provided as Docker containers [29] in order to represent a realistic behaviour as opposed to synthetic, simulated models. The configuration files of the ICT simulator integrate tightly with the rest of the software stack [30].

Co-Simulation

The multi-domain simulation for analysis can hardly be performed by one software tool alone. The setup of the last three sections describes three different, but intertwined, domains; each one warrants its own specific simulation software to yield realistic results [31]. In addition, specific models for power plants, wind parks, or independent market actors exist. These components are coordinated with the open-source co-simulation framework *mosaik* [32] and can therefore easily be integrated in other simulation setups relying on *mosaik*.

Figure 1 shows the complete software stack. The bottom box, labelled *co-simulation*, provides the technical view of the different simulators. Each simulator offers models, as well as attributes on these models, which form a hierarchy: The address scheme `Simulator.Model.Attribute` allows for unambiguous identification of each individual attribute and to connect them. E.g., `ARL.Attacker-1.Actuator-1` can be connected to `PowerGrid.WindFarm-1.P-Feedin` to deliver setpoints from the adaptive attacker agent to a wind farm under its control; similarly, `PowerGrid.Sensor-1.Voltage`, connected to `ARL.Attacker-1.Sensor-1`, allows the agent to measure the effects of its actions in terms of voltage values. *mosaik* synchronizes all simulators with each other and provides a common simulation clock time, the *time step*; data is transmitted to a simulator when it is *stepped*, data from its models' attributes is queried afterwards.

III. Distributed Analysis: Communication & Control

To analyze this interconnected complex system, the core tool is the application of the Adversarial Resilience Learning (ARL) methodology. ARL defines in its pure form [33] two classes of agents: *attacker agents* and *defender agents*. An instance of every class operates on a model of a CPS, i.e., both agents operate on the same shared model. However, neither attacker nor defender know of each other: They gather data from the CPS through their sensors, which retrieve the current state of the system—as far as it is observable to the respective agent—, but do not explicitly track changes induced by another party.

This specific distinction makes sense for the power grid, as well as for many other CPSs: Whether a voltage irregularity is induced by a larger PV feed-in at the end of the branch (e.g., coming from a farm) or forms a part of an attack, is hardly distinguishable, but needs to be countered in any case. Stringently, we assume that the defender needs to counter a variety of effects for resilient operation, from fluctuation in renewable feed-in to accidents to actual attacks without differentiating between them as a rule-based system would do. Therefore, neither the overall system design nor the experimenter differentiates between different causes and effects, leaving the development of strategies, as well as countering the adaption of the attacker to the defender's capability to adapt (and vice versa). That both agents learn to counter each other's strategies, thus developing them further and further, is the core of the system-of-systems learning principle of ARL [34]. Consequently, we use the attacker not just to execute actual cyber-attacks, but to represent any potentially system-harming behavior. Thus, the attacker becomes a universal analysis tool.

Focusing on the attacker, we consider a group of attacking ARL agents that form a self-organizing MAS and a single

 ISBN: 978-1-61208-818-1

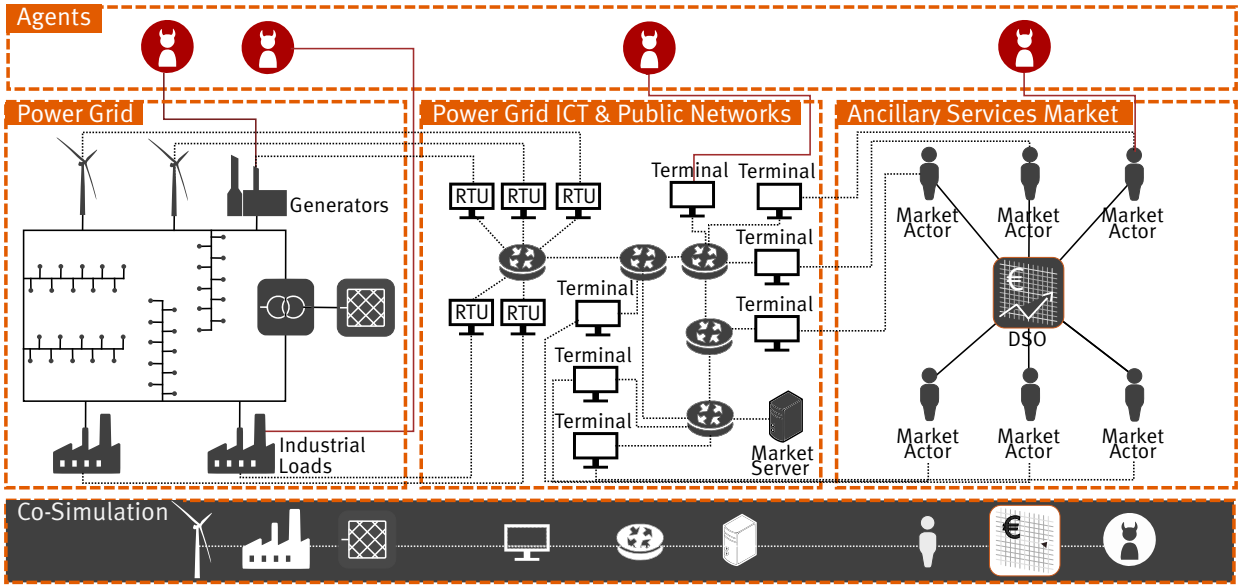

Figure 1. The Demonstrator's Co-Simulated Environments for Analysis

defender agent that represents the grid operator. All ARL agents use a modified Reinforcement Learning (RL) algorithm to explore a system that is initially unknown to them. In fact, ARL agents possess no domain-specific knowledge; their sensors and actuators contain only a description of the space for valid values. For the experimenter, these space types provide an easy way to describe types and boundaries for values; they can also be used as predicates to check whether a concrete value is a valid member of the given space. E.g., for a given value x, x is a member of the space $Discrete\{x\}$ iff:

$$Discrete\{n\} : x \in \mathbb{N}, 0 \leq x \leq n-1 \ . \tag{1}$$

Similarly, we can denote a box in $\mathbb{R}^n$ and check for a value x to be a member of it:

$$Box\{(l_1,\ldots,l_n),(h_1,\ldots,h_n)\} : x \in \mathbb{R}, \bigwedge_{i=1}^{n} l_i \leq x \leq h_i \ . \tag{2}$$

Other space types are *MultiDiscrete*, *MultiBinary*, or *Tuple*. Such a space description might represent the state of a tap-changer or the feed-in of a power plant in terms of a faction of its nominal output, but this logic is completely hidden from the agent. In fact, the domain logic is the responsibility of the experimenter. As the only way for RL agents to learn is to receive feedback, the experimenter has to derive a proper reward function that covers the relevant aspects of the CPS. The reward function bridges the otherwise separated concerns, i.e., the ARL perspective that eschews domain knowledge and the CPS domain. Thus, the ARL agents remain free of any domain-specific knowledge, as the reward function is a unit-less scalar and the objective can be learned.

Because of this feature, we describe the agents as being *polymorphic*. Drawing from the analogy in software engineering, the agents' interfaces are fixed and soundly described, but do not carry any model. That means, the space types assigned to the agents' sensors and actuators form a declaration, but no definition. The agents derive this definition—i.e., their model—through exploration. Hence, they are polymorphic. This means that an abstract definition of a CPS' interface in terms of the spaces outlined above is enough to have the ARL agents explore the systems; this constitutes a fundamental difference from many modelling and analysis tools that require implicit or explicit modelling of the target domain.

As part of this new research direction, we assume that MAS are a valid approach to analyze highly decentralized systems as depicted above: They inherently allow for a representation of local knowledge and rule sets, even learned one, such as a limited view on local grid state and local control options [35]. It has already been shown that a combination with cyber-physical energy system simulation is feasible and beneficial to analyze the distributed behavior of the system, even for socio-technical system views [36]. Thus, we use MAS to represent and explore the effect of cooperative malicious actors. In this case, cooperative means that the agents act cooperatively within their defined group of malicious or unplanned malicious, simply economically rational, agents. The attackers share a reward function, which can be as easy as the amount of money gained from the market, but also be complex and encompass aspects of all domains. In any case, the reward function remains transparent to each attacker and does not convey any domain knowledge to the agents, but is defined solely at the discretion of the experimenter.

In the presented concept, the overall MAS encompasses all three domains. Individual agents represent different actors in one of the domains. E.g., in a scenario, in which the attacker MAS controls three assets in the power grid, has one entry point to the ICT network, and appears with one bidder on the market, the MAS is comprised of five agents. An example for sensor and actuator mappings is presented in Table I.

 ISBN: 978-1-61208-818-1

TABLE I. EXEMPLARY ARL ATTACKER MAS THAT CAN PARTICIPATE IN A REACTIVE POWER MARKET

Agent	Asset	Sensors	Actuators
a_1	PV Unit	Voltage[1], Max. Active Power[2]	Active Power[2], Reactive Power[3]
a_2	EV Charger	Voltage[1], Active Power[2]	Active Power[2], Reactive Power[3]
a_3	Load	Voltage[1], Active Power[2], Reactive Power[3]	Active Power[2]
a_4	Market	Reactive Power Commitment (relative)[3]	Reactive Power Offer (relative)[3]
a_5	ICT	Interface Utilization[2]	Manipulate Sensor Value (Apply Noise)[3]

[1] $Box\{(0.85),(1.15)\}$, [2] $Box\{(0.0),(1.0)\}$, [3] $Box\{(-1.0),(1.0)\}$

In order to develop an overall strategy, the attackers need to coordinate among themselves without a central command-&-control instance. Snapshot algorithms [37] will be used to enable the agents to interchange their local sensor data to gain knowledge of the global state. In this case, global state means the entirety of all sensors that the ARL agents have access to. Learning agents that perform decision making based on shared knowledge can then learn optimal cooperative decision making based on that knowledge. With this research direction, we thrive for the development of a domain-encompassing coordination protocol to address this holistic approach to CPS analysis.

While malicious cooperation cannot be deduced directly from regulatory or observability loopholes, beneficial cooperative behavior is analyzed as well: the defender aims to stabilize the system and prevent malicious attacks.

In our research approach, we therefore combine these agent types to act in shared environments. Thus, we hope to identify ruleset, ICT, and market designs that minimize attack possibilities and stabilize the overall system. In future work, we will define and work out the resulting multi-layer attack coordination and defense framework.

IV. EXPERIMENT PROCESS

As Figure 2 illustrates, the overall experiment process incorporates four major steps: First, a domain independent description of the CPS and its interfaces is required. The definition of such a description is called *CPS Abstract Ontology* (CPS-AO) in the context of the presented research direction. The main purpose of the CPS-AO is the definition of network topological variables and the mapping of the ARL agents' sensors and actuators to entities in the environment. Additionally, the CPS-AO defines which variables can be changed during the experiments and the valid value ranges. Furthermore, the CPS-AO takes this topology information to build up experiments. For this purposes, CPS-AO employs techniques from the domain of Design of Experiments (DoE) [38] to select only configurations that provide the strongest significance. An example for a CPS-AO configuration file can be seen in Figure 3.

Furthermore, the CPS-AO serves as an input for the so-called *experiment generator* (CPS-EG). While the CPS-AO is a domain-independent and abstract description of the system, the CPS-EG instantiates the experiment descriptions for the actual simulation, assigns values to factors, and builds execution scripts.

All so generated concrete experiments are executed by an *experiment executor* (CPS-EE) in the target environment. This provides the actual interface between the agent structure and the simulation environment. In order to enable changes, the created intermediate results will be saved so that smaller changes are possible without having to go through the entire process again. During the execution of the experiments, the states of the simulation, as well as the actions of the agents and the market results are stored and thus made available for a later weak point analysis, which will be described in the following section.

V. POST-RUN ANALYSIS METHODOLOGY

The experimenter defines a set of invariants that describe the environment's overall health. After the executor has finished the simulation run and health invariants were falsified, a *post-mortem analysis* of the defeated system shall be conducted. This CPS Vulnerability Analyzer (CPS-VA) conducts targeted evaluation of the attacks across all domains, aiming to find the smallest chain of stringent actions that defeated this system, i.e. to identify the cross-domain attack-path or kill-chain. We assume that the ARL MAS, in its exploration, conducts a lot of negligible action before staging a successful attack. Hence, identification of the minimal kill-chain is a separate analysis task.

Another goal of the research project is the development of the CPS-VA, which primarily aims to understand the produced data. It operates on data from all nodes, i.e. data from sensors and actuators. The transactions on the market, as well as states from the ICT and the power grid are collected. Then, the CPS-VA is designed to apply different analysis techniques on this data to isolate the kill-chain. For understanding the path of the kill-chain, predictive models and techniques are often not the best choice [39]. Most of the time, causality methods are more promising. Mueller, Memory, and Bartrem [40] use causal discovery techniques to discover cyber kill-chains; therefore, data is presented as a Causal Bayesian Network (CBN). Finding the right methodology to explain the experiment's outcome is also part of the development of the CPS-VA and will start as soon as first datasets are ready.

From the ICT security point of view, the task of the CPS-VA is somehow similar to what threat detection tools are designed for, but so far they only focus on ICT related data. The behaviour of the ARL agents can be treated as Advanced Persistent Threats (APT). APTs can be described as sophisticated attack processes that are often strategically-motivated and profit-focused [41]. Standard industry solutions to detect APTs are so called Security Information and Event Management (SIEM) systems. Such a system collects data from a wide variety of security applications to detect suspicious traffic and behaviour in ICT systems. To make use of this information, SIEM systems use correlation rules [42] and rise alarm in case of a anomaly detection. The CPS-VA provides

 ISBN: 978-1-61208-818-1

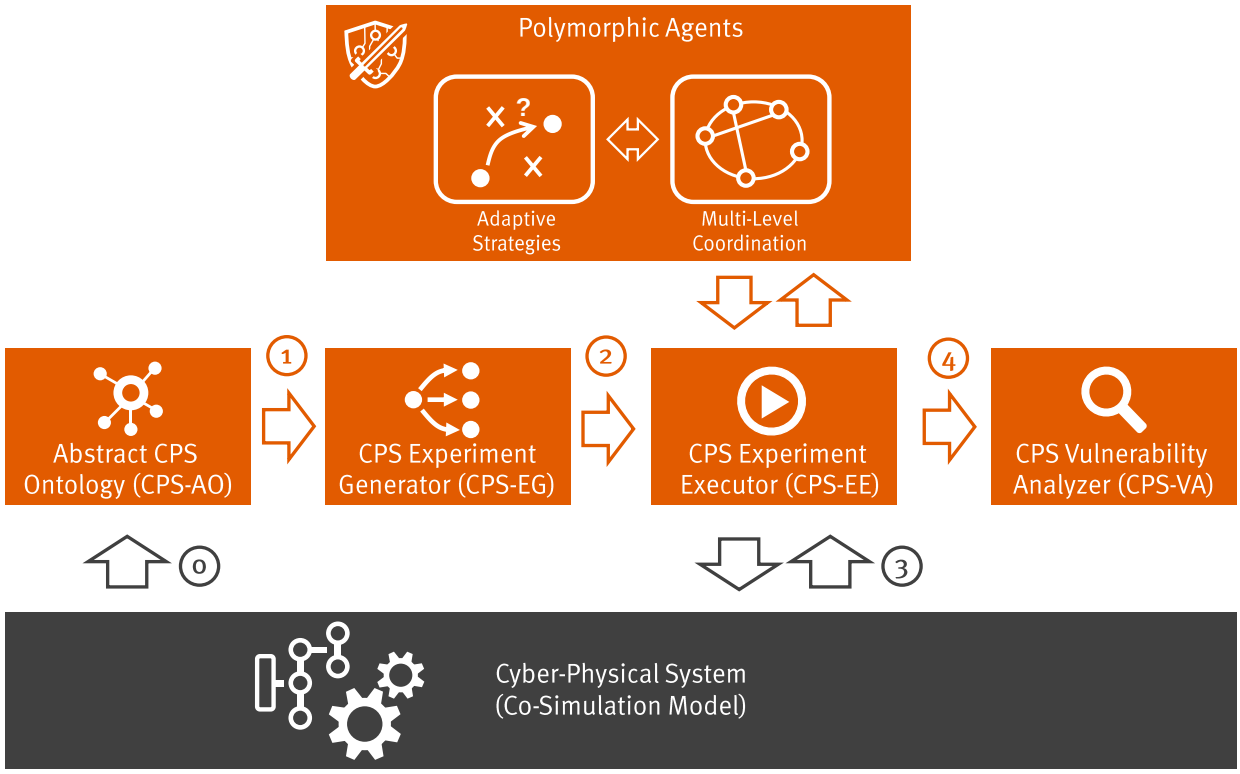

Figure 2. Experiment process of the presented research approach

```
!CPSAO
cps: !CPS # the system
  engine: # e.g. mosaik
  api: # how to instantiate
  sensors: # list of UIDs
  actuators: # list of UIDs
doe:
  runs:
  factors: # DoE inputs
  qualities: # DoE outputs
  strategy: # how to sample
  # e.g. Latin Hypercube

agents: # two or more
- !Agent
# if more than one option
# is present in the agent's
# definition, they will be
# considered for DoE
  name: # convenience
  sensors: # list of UIDs
  actuators: # list of UIDs
  strategies: # how to win
  rewards: #
- !Agent # same as above
```

Figure 3. A minimal example for a CPS-AO file. Additional parameters have been removed for the sake of brevity.

the opportunity to evaluate the idea of SIEM systems towards new applications.

First, a reasonable model from all domains is used in simulation to manually create simple correlation rules. This first step evaluates which information from the domains is necessary and how to create suitable correlation rules to generate basic knowledge for the next steps. Second, a much higher amount of relevant information for the SIEM is expected. In correlating different experiment runs from a variety of different scenarios—using, e.g., big data analytics [43]—, singular kill-chains can be derived and, thus, the respective rules be created. We expect that, starting with easy-to-observe critical states in the CPSs, an isolation path beginning on the affected components in the CPS can connect the critical states to market actors.

VI. Conclusion

Many CPS experience a broad addition of inputs, from self-driving capabilities over user inputs and IoT technologies to a broad market adoption in the case of power systems. The emergence of complex CPSs cannot be covered by traditional modelling and analysis techniques that can address only specific aspects of the overall system. In this paper, we proposed the concept for an application of distributed artificial intelligence as a self-adaptive analysis tool that is able to analyze the interdependencies between domains in CPS, covering the whole system. It eschews pre-configured domain knowledge, instead exploring the CPS domains for emergent risk situations and exploitable loopholes in codices, with a special focus on rational market actors that exploit the system while still following the rules of market.

In the future, we will demonstrate the feasibility of a cross-domain distributed analysis, documenting the experimentation system, the coordinating MAS-based exploration tool, as well as the analysis tool. With the latter, we aim to extract a reduced chain of actions leading to a cross-system exploitation, thereby isolating attack vectors and loopholes in codices. Furthermore,

 ISBN: 978-1-61208-818-1

we expect the use of polymorphic agents to lead to new insights in the field of RL. The ARL agent interaction with the ICT, which forms a central piece of the concept, will give new valuable insights of the ICT's critical role in modern CPSs. This will enhance research towards new security tools for modern critical infrastructures.

Currently, the implementation of this framework is not yet made available to the public; however, we expect this to happen in the coming weeks. We will then publish detailed comparisons to other approaches and make the respective data available for reproducing our results.

ACKNOWLEDGEMENTS

We would like to thank Sebastian Lehnhoff for his councel and valuable inputs. This work was funded by the German Federal Ministry of Education and Research through the project *PYRATE* (01IS19021A).

REFERENCES

[1] European Union, *Directive (EU) 2019/944 of the european parliament and of the council of 5 june 2019 on common rules for the internal market for electricity and amending directive 2012/27/EU*, [retrieved: Oct, 2020], 5 June 2019. [Online]. Available: https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32019L0944.

[2] E. M. Veith, B. Steinbach, and J. Windeln, "A lightweight distributed software agent for automatic demand—supply calculation in smart grids", *International Journal On Advances in Internet Technology*, vol. 7, no. 1, pp. 97–113, 2014.

[3] M. Ruppert, E. M. Veith, and B. Steinbach, "An evolutionary training algorithm for artificial neural networks with dynamic offspring spread and implicit gradient information", in *The Sixth International Conference on Emerging Network Intelligence (EMERGING 2014)*, International Academy, Research, and Industry Association (IARIA), IARIA XPS Press, 2014, pp. 18–21.

[4] E. M. Veith, *Universal Smart Grid Agent for Distributed Power Generation Management*. Berlin, Germany: Logos Verlag Berlin GmbH, Oct. 2017.

[5] O. Hanseth and C. Ciborra, *Risk, complexity and ICT*. Cheltenham, UK: Edward Elgar Publishing, 2007.

[6] D. Sculley *et al.*, "Machine learning: The high interest credit card of technical debt", *SE4ML: Software Engineering for Machine Learning (NIPS 2014 Workshop)*, pp. 1–9, 2014.

[7] D. U. Case, "Analysis of the cyber attack on the ukrainian power grid", *Electricity Information Sharing and Analysis Center (E-ISAC)*, 2016.

[8] J. Styczynski and N. Beach-Westmoreland, "When the lights went out: Ukraine cybersecurity threat briefing", *Booz Allen Hamilton*, vol. 12, p. 20, 2016.

[9] Reuters, *Ukrainian banks, electricity firm hit by fresh cyber attack*, Jun. 2017.

[10] A. Crawley, "Hiring hackers", *Network Security*, no. 9, pp. 13–15, Sep. 2016, ISSN: 13534858.

[11] K. Pei, Y. Cao, J. Yang, and S. Jana, "DeepXplore: Automated whitebox testing of deep learning systems", 2017, [retrieved: Oct, 2020]. arXiv: 1705.06640. [Online]. Available: http://arxiv.org/abs/1705.06640.

[12] T. Gehr, M. Mirman, D. Drachsler-Cohen, P. Tsankov, S. Chaudhuri, and M. Vechev, "AI^2: Safety and robustness certification of neural networks with abstract interpretation", in *2018 IEEE Symposium on Security and Privacy (SP)*, IEEE, 2018, pp. 1–18.

[13] S. Yaghoubi and G. Fainekos, "Gray-box adversarial testing for control systems with machine learning components", vol. 1, no. 1, pp. 179–184, 2019. arXiv: arXiv:1812.11958v1.

[14] A. Teixeira, S. Amin, H. Sandberg, K. H. Johansson, and S. S. Sastry, "Cyber security analysis of state estimators in electric power systems", *Proceedings of the IEEE Conference on Decision and Control*, pp. 5991–5998, 2010, ISSN: 01912216.

[15] H. Sandberg, A. Teixeira, and K. H. Johansson, "On security indices for state estimators in power networks", in *First Workshop on Secure Control Systems (SCS), Stockholm, 2010*, 2010, pp. 1–6.

[16] Y. Liu, P. Ning, and M. K. Reiter, "False data injection attacks against state estimation in electric power grids", *ACM Transactions on Information and System Security (TISSEC)*, vol. 14, no. 1, p. 13, 2011.

[17] S. Gao, L. Xie, A. Solar-Lezama, D. Serpanos, and H. Shrobe, "Automated vulnerability analysis of ac state estimation under constrained false data injection in electric power systems", in *Proceedings of the IEEE Conference on Decision and Control*, vol. 54, IEEE, 2015, pp. 2613–2620, ISBN: 9781479978861.

[18] L. Hu, Z. Wang, Q.-L. Han, and X. Liu, "State estimation under false data injection attacks: Security analysis and system protection", *Automatica*, vol. 87, pp. 176–183, 2018.

[19] P. Ju and X. Lin, "Adversarial attacks to distributed voltage control in power distribution networks with DERs", in *Proceedings of the Ninth International Conference on Future Energy Systems*, ACM, 2018, pp. 291–302, ISBN: 9781450357678.

[20] L. Hirth and I. Schlecht, "Market-based redispatch in zonal electricity markets", *SSRN Electronic Journal*, no. 055, pp. 1–26, 2018.

[21] C. Konstantinidis and G. Strbac, "Empirics of intraday and real-time markets in Europe: Great Britain", DIW – Deutsches Institut für Wirtschaftsforschung, Berlin, Germany, Tech. Rep., 2015, p. 21.

[22] E. M. Veith, L. Fischer, M. Tröschel, and A. Nieße, "Analyzing cyber-physical systems from the perspective of artificial intelligence", in *Proceedings of the 2019 International Conference on Artificial Intelligence, Robotics and Control*, ser. AIRC '19, Cairo, Egypt: Association for Computing Machinery, 2019, pp. 85–95, ISBN: 9781450376716.

[23] K. Rudion, A. Orths, Z. A. Styczynski, and K. Strunz, "Design of benchmark of medium voltage distribution network for investigation of dg integration", in *2006 IEEE Power Engineering Society General Meeting*, IEEE, 2006, pp. 6–12.

[24] CIGRE Task Force C6.04.02, *Benchmark Systems for Network Integration of Renewable and Distributed Energy Resources*. 2014.

[25] L. Hofmann and M. Sonnenschein, "Smart Nord—final report", *Hartmann GmbH*, 2015.

[26] C. Weinhardt *et al.*, "How far along are local energy markets in the DACH+ region?: A comparative market engineering approach", in *Proceedings of the Tenth ACM International Conference on Future Energy Systems, e-Energy 2019, Phoenix, AZ, USA, June 25-28, 2019*, ACM, 2019, pp. 544–549.

[27] F. Lilliu, M. Vinyals, R. Denysiuk, and D. R. Recupero, "A novel payment scheme for trading renewable energy in smart grid", in *Proceedings of the Tenth ACM International Conference on Future Energy Systems, e-Energy 2019, Phoenix, AZ, USA, June 25-28, 2019*, ACM, 2019, pp. 111–115.

[28] S. C. Chau, J. Xu, W. Bow, and K. M. Elbassioni, "Peer-to-peer energy sharing: Effective cost-sharing mechanisms and social efficiency", in *Proceedings of the Tenth ACM International Conference on Future Energy Systems, e-Energy 2019, Phoenix, AZ, USA, June 25-28, 2019*, ACM, 2019, pp. 215–225.

[29] The Docker developers, *Docker website*, [retrieved: Oct, 2020]. [Online]. Available: https://www.docker.com/.

 ISBN: 978-1-61208-818-1

[30] T. Woltjen, G. Gritzan, P. Kathmann, and R. Sethmann, "Simulationsumgebung für IKT-Netze zur Cyber-Abwehr", in *Tagungsband AALE 2020*, VDE Verlag, 2020, pp. 233–239, ISBN: 978-3-8007-5180-5.

[31] S. Balduin, M. Tröschel, and S. Lehnhoff, "Towards domain-specific surrogate models for smart grid co-simulation", *Energy Informatics*, vol. 2, no. 1, p. 27, 2019.

[32] The mosaik Developers, *Mosaik website*, [retrieved: Oct, 2020]. [Online]. Available: https://mosaik.offis.de/.

[33] L. Fischer, J.-M. Memmen, E. M. Veith, and M. Tröschel, "Adversarial resilience learning—towards systemic vulnerability analysis for large and complex systems", in *The Ninth International Conference on Smart Grids, Green Communications and IT Energy-aware Technologies (ENERGY 2019)*, vol. 9, 2019, pp. 24–32.

[34] E. M. Veith, N. Wenninghoff, and E. Frost, *The adversarial resilience learning architecture for ai-based modelling, exploration, and operation of complex cyber-physical systems*, 2020. arXiv: 2005.13601 [cs.AI].

[35] Y. Shoham and K. Leyton-Brown, *Multiagent Systems - Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge, MA, USA: Cambridge University Press, 2009, ISBN: 978-0-521-89943-7.

[36] I. Praca, H. Morais, C. Ramos, Z. Vale, and H. Khodr, "Multi-agent electricity market simulation with dynamic strategies & virtual power producers", in *2008 IEEE Power & Energy Society general meeting*, Piscataway, NJ: IEEE, 2008, pp. 1–8, ISBN: 978-1-4244-1905-0.

[37] K. M. Chandy and L. Lamport, "Distributed Snapshots: Determining Global States of Distributed Systems", *ACM Transactions on Computer Systems (TOCS)*, vol. 3, no. 1, pp. 63–75, 1985, ISSN: 07342071.

[38] J. P. Kleijnen, "Design and analysis of simulation experiments", in *International Workshop on Simulation*, Springer, 2015, pp. 3–22.

[39] G. Shmueli, "To Explain or to Predict?", *Statistical Science*, vol. 25, no. 3, pp. 289–310, Aug. 2010, [retrieved: Oct, 2020], ISSN: 0883-4237. arXiv: 1101.0891. [Online]. Available: http://projecteuclid.org/euclid.ss/1294167961.

[40] W. G. Mueller, A. Memory, and K. Bartrem, "Causal discovery of cyber attack phases", *Proceedings - 18th IEEE International Conference on Machine Learning and Applications, ICMLA 2019*, pp. 1348–1352, 2019.

[41] A. Ahmad, J. Webb, K. C. Desouza, and J. Boorman, "Strategically-motivated advanced persistent threat: Definition, process, tactics and a disinformation model of counterattack", *Computers & Security*, vol. 86, pp. 402–418, Sep. 2019, ISSN: 01674048.

[42] A. Ambre and N. Shekokar, "Insider Threat Detection Using Log Analysis and Event Correlation", *Procedia Computer Science*, vol. 45, no. C, pp. 436–445, 2015, ISSN: 18770509.

[43] A. a. Cardenas, P. K. Manadhata, and S. P. Rajan, "Big data analytics for security", *IEEE Security & Privacy*, vol. 11, no. 6, pp. 74–76, Nov. 2013, ISSN: 1540-7993.