Digital-Twin-Driven Deception Platform: Vision and Way Forward

Sabah Suhail, *Queen's University Belfast, Belfast, BT7 1NN, U.K.*
Mubashar Iqbal, *University of Tartu, Tartu, 51009, Estonia*
Kieran McLaughlin, *Queen's University Belfast, Belfast, BT7 1NN, U.K.*

Digital twin (DT) technology provides new opportunities to enhance the robustness and resilience of critical infrastructure. In this article, we discuss the potential of DTs as a proactive security enabler and present a vision for using DTs as a deception platform to thwart cyberattacks on critical national infrastructure (CNI). We propose a generic DT-driven deception-based solution called securing cyberphysical systems through DT-driven deception (INCEPTION), which can serve as a research-based DT-driven deception platform for CNI. INCEPTION is intended to provide a set of foundational components, establishing a baseline that supports further developmental work within the realm of cybersecurity deception. Furthermore, we discuss the way forward by identifying research challenges in DT-driven deception solutions.

The evolving landscape of cyberattacks related to critical national infrastructure (CNI), specifically targeting energy and utility sectors, has led to repercussions affecting social, economic, and business sectors.[1,2] Seminal examples of cyberattacks on CNI include breaches in water systems[3] and disruptions in gas pipelines.[4] Adopting proactive security mechanisms to enhance security by identifying, detecting, and responding to ongoing attacks in cyberphysical systems (CPSs) is imperative. Proactive security solutions, such as honeypots, deal with thwarting potential attacks and identifying future attack vectors. Honeypots intentionally entice cyberattackers, drawing them into compromise and attack scenarios to identify and study their tactics.[5] This way, actionable intelligence can be gathered on the attackers.[5]

However, in the wake of sophisticated attacks and advancements in myriad attack vectors, traditional honeypots cannot be rendered as the only line of defense against nontraditional intelligent and persistent attacks. This limitation is due to the absence of automation, scalability, and target specialization rather than generalization in the traditional honeypots.[6] Moreover, most existing honeypot, honeynet, or decoy systems in CPSs do not have high interaction and lack a fully realistic implementation because of the cost associated with setting them up and managing them.[5]

To this end, deception technology that leverages a range of available methods, such as honeypots, decoys, and breadcrumbs, can address the limitations of traditional proactive security solutions.[6] However, the management of deception within an operational CPS can be a challenge for several reasons: the underlying subsystems are dynamic and complex, evasion techniques might outweigh the deception, and the requirement for testing, updating, and redeployment of deception elements through extensive testbeds and experimentation platforms could be expensive.[7] In this regard, digital twins (DTs) emerge as a valuable strategy. A DT is a virtual (digital) representation of its physical counterpart along with its behavior and interconnections.[1] DTs provide a platform to design and implement a dynamic and scalable deception solution.

LEVERAGING DTs AS A CYBERDECEPTION PLATFORM

The concept of using DTs as a honeypot was first introduced by Eckhart and Ekelhart[1] to yield realism and resemblance to CPS infrastructure. Similarly, Eckhart et al.[10] proposed using DTs for cyberdeception.

Digital Object Identifier 10.1109/MIC.2024.3406188
Date of publication 7 June 2024; date of current version 2 August 2024.

TABLE 1. Mapping cyberdeception design and operational requirements with generic characteristics of DTs and INCEPTION—a DT-driven deception platform.*

Cyberdeception Requirements	Generic Characteristics of DTs	DT-Driven Deception Platform (INCEPTION)
Virtual/physical deception environment with level of interaction[5]	Low-, medium-, and high-fidelity simulation mode as well as testbeds	Twisted fidelity deceptive twin (refers to scaling down the level of representation and contextualization)
Intertwined realism and misrepresentation[8]	Clone counterpart through simulation mode	Instantiation (refers to creating versions of DTs based on a reference DT), twisted fidelity, and delusional attack surface
Scalability[5]	Simulation mode	Instantiation module to (re)generate instance(s)
Flexibility[7]	N/A	Deceptive twin scope
Proactive probing throughout cyber kill chain stages[8]	Simulation mode	Modus operandi of an adversary (IoCs, TTPs, and IoAs)
Reproducible/repeatable experiments to investigate adversary artifacts[9]	Simulation mode to run and rerun scenarios	Instantiation
Adversary continuous engagement[8]	N/A	Cyberdeception artifacts (Delusional attack surface, service/utility, data generation)
Updating and redeployment of deception elements[9]	N/A	Adaptive deception (state estimation and revise strategy), instantiation
Simulated services[5]	Simulation mode	Simulation mode with twisted fidelity through instantiation

*IoA: indicator of attack; IoC: indicator of compromise; N/A: not applicable; TTP: tactics, techniques, and procedures.

However, these works merely scratch the surface of harnessing DTs as honeypots or for cyberdeception, offering preliminary insights rather than comprehensive solutions. Leveraging DTs as a deception platform is possible when a DT is capable of providing a virtual replica that can simulate the system's structure, behavior, or context sufficiently to divert an attacker from the real system. DT approaches optimized for instantiation, interaction, and communication could become game changers for deception with multifaceted objectives, e.g., to lure the attackers to study their behaviors, identify risks and attacks, and mitigate future attacks. Thus, DTs can provide a favorable platform for realizing deceptive technologies in an agile, engaging, efficient, and dynamic way. Table 1 (the first and second columns) discusses the DT characteristics that make them valuable to be leveraged as a proactive security mechanism for CPS.

DTs: An Attacker's New Playground

Leveraging DTs as a proactive security mechanism could make them an attacker's new playground. Attackers may exploit the deep knowledge about physical devices accessible through DTs to steer the CPS into an insecure state.[1,10] Therefore, when employing DTs as a security enabler, it is important to consider the possibility of DT compromise.

Although the inbuilt characteristics of DTs somehow fulfill the requirements of cyberdeception (as the second column of Table 1 shows), it gives rise to other challenges. On the one hand, DTs as cyberdeception platforms are required to provide plausible-looking yet misleading information to deceive adversaries,[1] i.e., a sufficient level of interaction or medium- to high-fidelity deception. On the other hand, DTs risk revealing substantial information about the system to attackers. Considering both contradictory requirements, if DTs are used for cyberdeception, attackers will gain significant knowledge about the system, and, if DTs are modified, they cannot serve as a cyberdeception solution. For instance, attackers may exploit a deception layer to compromise the system and then execute other operations, such as jeopardizing other devices.[5] This problem was highlighted by Eckhart and Ekelhart[1] as follows: "How can existing DTs be adapted cost-effectively to convincingly simulate plant operations, ensuring that attackers, misled by these honeypots, cannot acquire valuable insights into the actual systems?"

To address such challenges, we envision a DT-driven deception solution called securing CPSs through DT-driven deception (INCEPTION).

INCEPTION

Figure 1 represents INCEPTION, a research-based DT-driven deception platform for CNI. A research-based deception platform collects and analyzes cyberattack data, helping develop stronger defenses and understand evolving tactics.[5] Table 1 shows the mapping of INCEPTION modules (the third column) with deception requirements. Harnessing DT as a proactive security mechanism entails four strands:

1) Designing a tailored deception environment strategically positioned to divert attackers from the real system.
2) Achieving high interaction with attackers in the deceptive environment by designing and deploying a deceptive twin.
3) Through deception orchestrator (DICE), analyzing attackers' engagement with the deceptive DT, identifying the vulnerabilities they target for exploitation, and anticipating potential future strategies.
4) Analyzing results to plan further actions regarding cybersecurity strategies.

Setting the Trap

The initial step in deploying a DT-driven deception solution is to steer attackers toward the deceptive environment. This is achieved by strategically placing decoys or breadcrumbs (baiting) within the real network environment. These decoys may include simulated vulnerabilities, fabricated credentials, or misleading links. The primary goal is diverting the attackers from the real system to a DT-driven deception environment.

Deceptive Twin Layer

Generally, the design and deployment of a DT-driven deception is steered by its scope. The scope defines which deception layer (network, system, data, or application) counters which attack.[7]

To realize DT-driven deception twins, low-fidelity twins can be instantiated from a reference DT. The reference DT functions as a comprehensive blueprint, replicating the physical system precisely and accurately. Given that the low-fidelity DT may lead to information leakage, a twisted fidelity module scales down the level of representation and contextualization, creating a simplified yet strategically deceptive environment. Thus, a twisted fidelity DT engages the attacker by providing a plausible-looking yet misleading virtual replica of the physical system. Moreover, deceptive

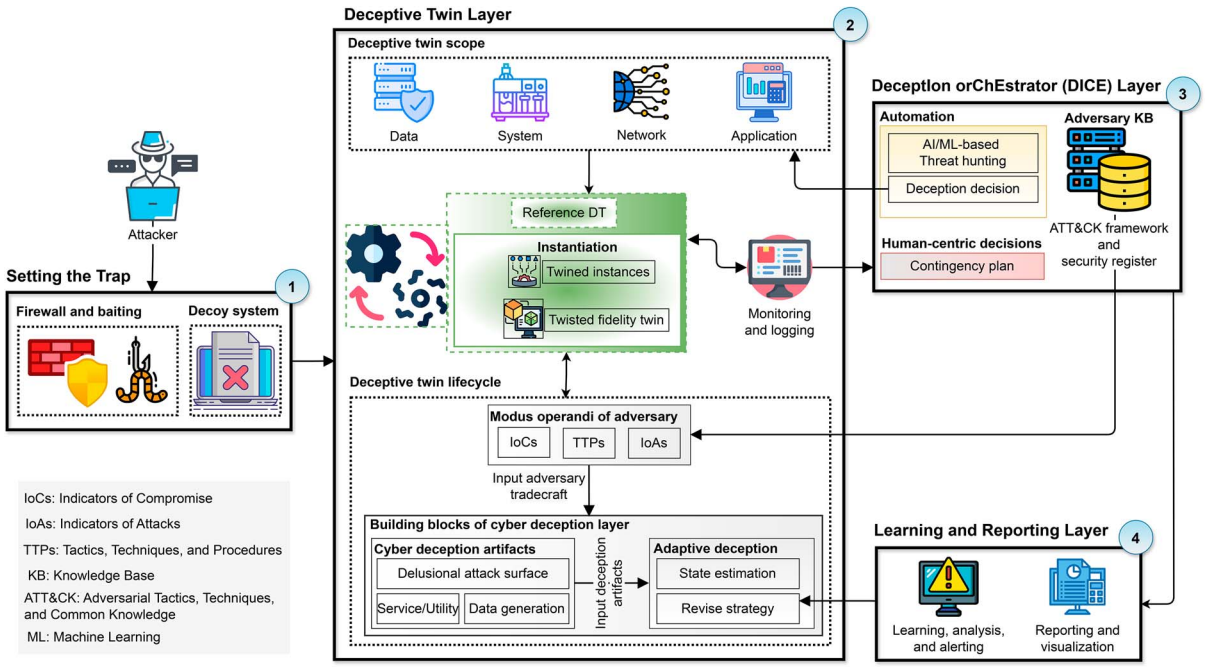

FIGURE 1. INCEPTION: A DT-driven deception solution. ① Setting a trap to divert attackers from the real system. ② Designing a deceptive twin that realizes a twisted fidelity. ③ Analyzing the attacker's engagement with the deceptive DT and identifying vulnerabilities. ④ Learning and reporting results for further actions.

twin instance versions can be managed and maintained to serve threefold purposes: 1) ensuring each instance is consistently updated to address evolving cyberthreats, 2) providing a platform for researchers to analyze attack methodologies and defensive strategies, and 3) facilitating customization and scalability to meet specific organizational requirements.

The deceptive twin lifecycle module provides deception operation planning and development. First, it involves gathering the attacker's modus operandi, which comprises indicators of compromise (IoCs); tactics, techniques, and procedures (TTPs), and indicators of attacks (IoAs) to spot the different forms of malicious activities. Second, cyberdeception artifacts provide a set of utilities and services composed of true and fabricated data and intentional delusional attack surfaces. Third, the adaptive deception strategy changes in response to the attacker's actions.

Following the deployment categorization defined by Han et al.[9] INCEPTION adopts a mixture of *added to* and *set in front of* the system. Such multideception placement and deployment protect the system from insiders and outsiders. During a multistage attack, the deployment strategy, i.e., added to the system, takes care of insiders, whereas positioning in front of the system targets outsiders. Given the dynamic nature of cyberthreats, the attacker's state estimation is required to formulate a deception layer.[8] Furthermore, the deception strategy needs to be revised on AI/machine learning (ML)-based solutions to engage attackers actively with a different deception approach. The attacker's state and action are logged and can be accessed by a deceptive twin lifecycle module.

DICE Layer

DICE is used to study the attacker's behavior within the cybersecurity landscape. DICE is based on passive and proactive probing, such as 1) data obtained through an adversary knowledge base (KB) and logged data and 2) AI/ML-based threat hunting to provide timely indications of presumably ongoing attacks. DICE consists of two submodules related to automation and humancentric decisions.

Context-aware security automation is essential due to the dynamic threat landscape. Integrating AI/ML-based threat hunting can significantly reduce the resources and effort of the cybersecurity assessment team.[11] Threat hunting is an iterative process of searching adversarial states and actions throughout the cyber kill chain (CKC) stages.[11] CKC represents the adversary's steps in a cyberattack, including reconnaissance, intrusion, lateral movement, privilege escalation, exploitation, and exfiltration, with the sequence varying, but effective adversary engagement is paramount in combating attacks. INCEPTION can support the hunting process to identify potential threats and attacks, such as reconnaissance and intrusion, during early CKC stages. The INCEPTION platform can be devised as an initial step based on the adversarial state information acquired from available crowdsourced attack data sources, such as MITRE ATT&CK for industrial control systems (ICSs). In practice, the INCEPTION platform may improve over time[8] by using the system's telemetry data (e.g., system logs, app logs, network traces) and can provide novel attack vectors to the adversary KB.

The deception decision module serves two purposes. *First*, it should establish a reactive cyberdefense posture. To do so, the deception decision module operates based on the threat-hunting outcome, triggering the cyberdeception layer module for insiders and outsiders. Furthermore, it must take decisions on deception operation continuation, adaptation, and termination depending on the context. *Second*, it should establish an active cyberdefense posture. With a research-based objective, deceptions must be deployed against critical data or services to study and anticipate the attacker's action. To do so, the deception decision module relies on the adversary KB and triggers the cyberdeception layer module for respective deception deployment.

Considering human–machine collaboration scenarios in CNI, it is essential to define fail-safe procedures and emergency provisions [e.g., isolate or shutdown (sub)systems] for human, physical, and environmental security and safety.[12] In ICSs, safety instrumented systems are designed for detecting violations against predefined safe state conditions and triggering isolation or shutdown of the device(s)/process(es).[12] However, completely automating security-critical functions might not be a feasible solution, and no amount of security control parameters[12] can protect infrastructure against persistent human adversaries. The automated system may be unable to process unforeseen circumstances due to biases, configuration settings, or resource constraints and, thus, requires a human in or on the loop.

For example, during troubleshooting, operators are misdirected by the human–machine interface showing closed valves when they are physically open, revealing the limitations of automated security responses. Depending on the event's context and severity of impact, DICE supports reliance on human knowledge (plant operators and security professionals), which can be supplemented by automated security solutions to identify root causes. The contingency plan can be based on the standard policies and procedures by the National Institute of Standards and Technology[12] to

allow graceful shutdown or isolate compromised subsystems and restore disrupted services upon recovery. Thus, instead of entirely relying on automated systems or humans, INCEPTION proposes combining both to follow timely countermeasures.

Learning and Reporting Layer

The DT-driven deception infrastructure can evolve based on a comprehensive understanding of how the attackers operate in response to deception measures. Including a visualization interface illustrates the effectiveness of deception measures taken in steps 2 and 3 and aids in making well-informed decisions regarding cybersecurity strategies. Furthermore, alerts are triggered based on identified suspicious activities, enabling rapid responses to cyberthreats. Insights from this analysis are then fed back into the deception module (step 2) to refine the deception strategy for countering the attacker's methods.

Overall, the INCEPTION platform contributes to the adversary KB in two ways: 1) the modus operandi of the adversary, which evolves as the attacker moves deeper in the deceptive twin, and 2) threat hunting through the DICE module, which provides early indications or response to security alerts if there are ongoing attacks on the deceptive twin.

USE CASE: ENHANCING WATER SYSTEM SECURITY THROUGH INCEPTION

Water treatment and distribution systems are CNIs that include reservoirs, treatment plants, pumping stations, and a network of pipes that deliver clean water to consumers. As reliance on digitization and interconnected systems grows, water utilities must be increasingly prepared for cyberthreats.

Cyberattacks on Water Systems

In a theoretical scenario, an attacker might gain access to a water system's network, setting the stage for a series of malicious actions. For instance, by exploiting supervisory control and data acquisition system vulnerabilities, the attacker might manipulate water treatment parameters, disable critical alarms, inject false data into sensors, or disrupt water distribution by altering valve operations or tampering with water pump pressure. Such security incidents could pose significant public health risks. To date, numerous cyberattacks have been reported on water systems. For instance, in the United States, water systems have been targeted in cyberattacks that resulted in data breaches and unauthorized modifications.[3] Such occurrences highlight the pressing necessity for robust cybersecurity measures to safeguard sensitive information and uphold the uninterrupted provision of essential services.

INCEPTION in Action

With the existing honeypots for water systems as surveyed by Franco et al.,[5] the following discussion answers the following question (identified by Han et al.[9]): How can a DT-driven deception platform be effectively integrated with existing systems, addressing their drawbacks, such as lack of reproducible experiments to test deception and updating and redeployment of deception elements?

In response to the growing cybersecurity threats to CNI, implementing the INCEPTION solution (Figure 2) can fortify the water system defense. Starting with setting the trap (Figure 1: step 1), decoys and baiting (e.g., a firewall, an open port strategy, or leaked credentials) are strategically placed within the water system's network to lure attackers into a controlled DT-driven deceptive environment.

Continuing with the design and deployment of deceptive twin (Figure 1: step 2), a twisted fidelity twin mimicking the water control system is instantiated from a reference DT of the water system. Under practical settings, the deceptive environment is designed based on available resources and the assets to be protected. While engaging with INCEPTION, the attacker may attempt one or several of the previously mentioned cyberattacks depending on their motives and objectives. Throughout the attacker engagement period with the deceptive twin, attacker behavioral analytics and profiling activities are performed (Figure 1: step 3). Finally, the learning and reporting phase (Figure 1: step 4) analyzes results to identify actionable measures to enhance the security of the real water system.

Expanding upon existing virtual honeypots for water systems, the integration of INCEPTION enables the instantiation of variable fidelity DT instances. This integration facilitates the reproduction and repetition of experiments investigating attack artifacts on the water system. Moreover, instantiation enhances scalability compared to physical testbeds, which are often time-consuming and cost-intensive to set up and maintain and have limited scalability and flexibility. Given the spectrum of interaction levels observed in existing honeypots as surveyed by Franco et al.,[5] which range from low to medium or high depending on available resources and intended purpose, the variable twisted fidelity proposed by INCEPTION can accommodate

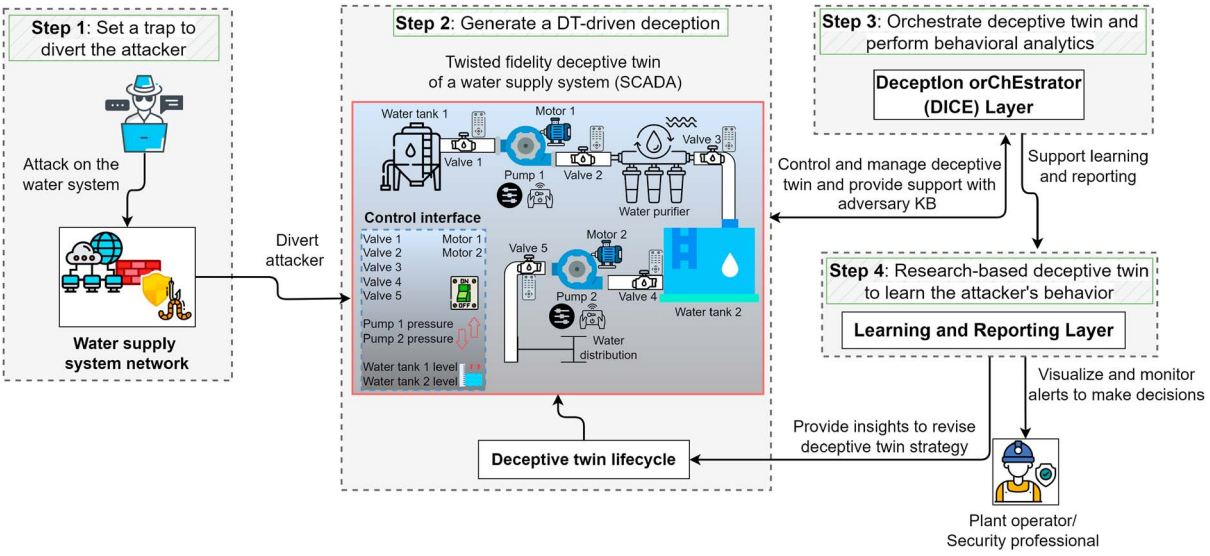

FIGURE 2. Integrating INCEPTION (Figure 1) to enhance water supply system security. Step 1: Set a trap to divert the attacker from the real water supply system. Step 2: Generate a twisted fidelity deceptive twin. For illustration purposes, the primary components depicted are water tanks, valve controls, pump stations, water flow, and pressure systems. Step 3: Use the DICE layer to identify threats and vulnerabilities. Step 4: Monitor and report the attacker's movements to plant operators or security professionals.

these diverse requirements. This adaptability is advantageous, as high interaction capabilities are crucial for identifying complex attacks, while low interaction can yield valuable insights into attack origin and brute-force attempts.[5] Furthermore, the ability to adapt between low- and high-fidelity levels makes them viable options for smaller organizations or resource-constrained ones.

Amidst evolving cyber tradecraft within CNI, INCEPTION can reinforce CNI defenses, provided the following challenges are addressed.

CHALLENGES AND WAY FORWARD

We cover a spectrum of considerations, from technological hurdles to strategic deployment challenges, aiming to highlight the practicality, efficiency, and future direction of leveraging DT technology for cybersecurity measures (Figure 3). INCEPTION addresses some of these challenges, while others warrant further investigation and research.

Cost–Benefit Analysis

Deploying additional DT-driven deception layers may incur capital and operational expenditures. Therefore, the overarching question is this: Would the return on investment be of significant value to justify the

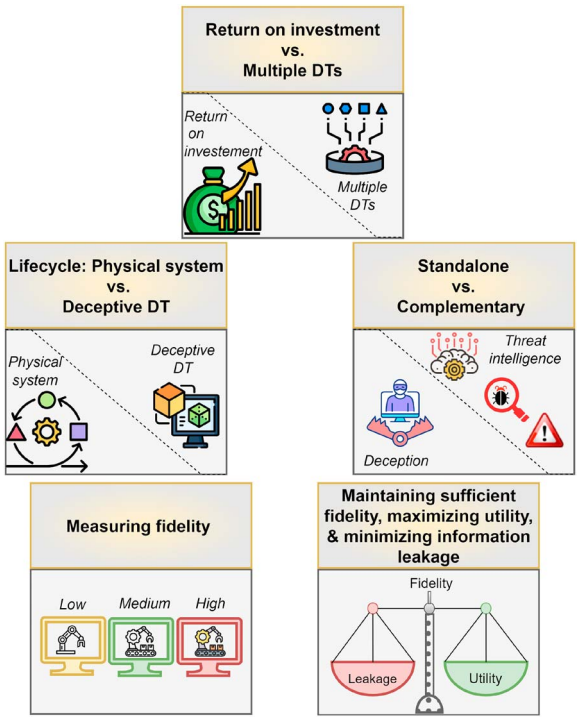

FIGURE 3. DT-driven deception solutions: challenges and open-ended questions.

rationale of creating (multiple) variable twisted fidelity DTs? Considering the balance between security and operational performance, such a solution is more suitable for realizing the security-enhancing objectives of the CNI. This notion holds significant value where state-aligned actors target CNI, such as attacks on water and wastewater sectors[3] or gas pipelines.[4]

Nevertheless, based on the cost–benefit analysis of the underlying application areas, the expenditures, including monetary, computation, communication, storage, and maintenance, could be reduced by limiting the simulated services or providing only essential services. For instance, the INCEPTION solution (Figure 1: step 2) can adapt its modules, including instantiation, deceptive twin scope, and cyberdeception artifacts, based on available resources, severity, and potential impact of cyberattacks. While mapping an entire infrastructure to a DT for cybersecurity analysis is often impractical,[13] instead of a single, monolithic DT of the water system (including all components), as step 2 in Figure 2 shows, a cost-efficient approach to DT development includes creating DT instance(s) for individual critical components, such as networks, systems, applications, or combinations thereof, to serve as a deception platform. Nevertheless, organizations may still need to address challenges related to cost management and the retrofitting of legacy systems when implementing such strategies.

Deceptive Twins to Evolving Cyberthreats

The lifecycle of DTs, i.e., the design and operation, needs to evolve with the lifecycle and complexity of the CPS. Such requirements are especially important for a high-fidelity security-enhancing DT due to state synchronization. However, the question is how to control a deceptive twin with the evolution of the physical system (and its twinned instances) and with the attacker's movements. To address this requirement, cyberdeception can integrate moving target defense (MTD) strategies that continuously transform the deception layer based on the attacker's cognitive bias, i.e., the attacker's plan of action due to misinformation, moving decisions, and associated configurations.[8] In this regard, using adaptive or self-configuring cognitive DTs offers a viable strategy.[14] When employed as deceptive DTs, such cognitive DTs can autonomously adjust their configurations, providing continuous safeguarding against evolving cyberthreats while prolonging engagement with the attacker.

Stand-Alone or Complementary?

Another deciding factor for deploying a deception platform is to consider the option of deception as a stand-alone solution or complementary solution to other defense approaches, as suggested by Han et al.[9] A DT-driven deception may extend other services or solutions, particularly when twisted fidelity DTs are needed to instantiate variable DT instances to support testing system vulnerabilities, training personnel, and conducting diagnostics. Furthermore, insights from deception-based solutions have the potential to significantly enhance an organization's cybersecurity posture by integrating the organization's threat intelligence database with IoCs and TTPs identified from the deceptive engagements. For instance, steps 3 and 4 in Figure 1 can play a significant role in detecting, deterring, and responding to cyberthreats in water systems.

Twisted Fidelity

The deceptive twin must have sufficient fidelity to delay an attack in progress while engaging the attacker. However, for insider threats and attackers, the stringent requirement of sufficiently high-fidelity DTs may contradict the objective of a variable (low) fidelity deceptive twin. The fidelity of a DT is based on several factors, including 1) the underlying objectives of leveraging DTs, 2) the degree of similarity between a physical asset and its virtual counterpart, 3) digital–physical synchronization frequency of data, 4) the tradeoff between generalization and contextualization, 5) resource constraints, and 6) DT modeling complexity and accuracy.[2] Such characteristics associated with measuring the fidelity of a DT raise this question: For DT-driven deception, how difficult is it to measure fidelity?

While fidelity requirements vary between systems, DT-driven deception can utilize fidelity by offering utilities and services with both authentic and fabricated data, such as sensors that relay fake and real data streams. High fidelity within a deception environment, characterized by realism and resemblance to the CPS infrastructure, enhances the utility by facilitating continuous engagement with attackers. However, such benefits come at the expense of information leakage, exposing the functioning of the physical system,[2] requiring additional resources, and escalating the complexity of the honeypot. Thus, investigating the tradeoff between the level of fidelity versus the utility of the deceptive DT and its effect on infrastructure-related information leakage is an open research question.

In addition to addressing the aforementioned challenges, investigating the feasibility and efficacy of DT-driven deception solutions in diverse CNI settings offers a promising direction for future research. For instance, further exploration is recommended toward

achieving twisted fidelity to address the problem of information leakage through DTs.

CONCLUSION

This article envisions INCEPTION, a research-based, DT-driven deception approach. There is an emerging opportunity to develop DT-driven deception solutions that can combine proactive cyberdefense techniques, including cyberdeception and MTD, for learning the attacker's goals. We anticipate that the INCEPTION approach establishes a baseline that can inspire further investigation by the research community to fully realize a toolchain for DT-driven deception.

ACKNOWLEDGMENTS

This work was supported by the North-South Research Programme, which is delivered by the Higher Education Authority on behalf of the Department of Further and Higher Education, Research, Innovation and Science and the Shared Island Unit at the Department of the Taoiseach.

REFERENCES

1. M. Eckhart and A. Ekelhart, *Digital Twins for Cyber-Physical Systems Security: State of the Art and Outlook*. Cham, Switzerland: Springer-Verlag, 2019, pp. 383–412.
2. S. Suhail et al., "The perils of leveraging evil digital twins as security-enhancing enablers," *Commun. ACM*, vol. 67, no. 1, pp. 39–42, 2023.
3. N. Tuptuk et al., "A systematic review of the state of cyber-security in water systems," *Water*, vol. 13, no. 1, 2021, Art. no. 81, doi: 10.3390/w13010081.
4. J. Beerman, D. Berent, Z. Falter, and S. Bhunia, "A review of colonial pipeline ransomware attack," in *Proc. IEEE/ACM 23rd Int. Symp. Cluster, Cloud Internet Comput. Workshops (CCGridW)*, Bangalore, India, 2023, pp. 8–15, doi: 10.1109/CCGridW59191.2023.00017.
5. J. Franco, A. Aris, B. Canberk, and A. S. Uluagac, "A survey of honeypots and honeynets for Internet of things, industrial Internet of things, and cyber-physical systems," *IEEE Commun. Surveys Tuts.*, vol. 23, no. 4, pp. 2351–2383, Fourthquarter 2021, doi: 10.1109/COMST.2021.3106669.
6. W. Lance. "Goodbye, honeypots–hello, true deception technology." Security Magazine. Accessed: Jan. 31, 2024. [Online] Available: https://www.securitymagazine.com/articles/92903-goodbye-honeypots-hello-true-deception-technology
7. L. Zhang and V. Thing, "Three decades of deception techniques in active cyber defense—Retrospect and outlook," *Comput. Secur.*, vol. 106, Jul. 2021, Art. no. 102288.
8. C. Wang and Z. Lu, "Cyber deception: Overview and the road ahead," *IEEE Security Privacy*, vol. 16, no. 2, pp. 80–85, Mar./Apr. 2018.
9. X. Han et al., "Deception techniques in computer security: A research perspective," *ACM Comput. Surv.*, vol. 51, no. 4, pp. 1–36, 2018.
10. M. Eckhart et al., "Security-enhancing digital twins: Characteristics, indicators, and future perspectives," *IEEE Security Privacy*, vol. 21, no. 6, pp. 64–75, Nov./Dec. 2023, doi: 10.1109/MSEC.2023.3271225.
11. B. Nour, M. Pourzandi, and M. Debbabi, "A survey on threat hunting in enterprise networks," *IEEE Commun. Surveys Tuts.*, vol. 25, no. 4, pp. 2299–2324, Fourthquarter 2023, doi: 10.1109/COMST.2023.3299519.
12. K. Stouffer et al. "Guide to operational technology (OT) security." NIST. Accessed: Jan. 10, 2024. [Online]. Available: https://doi.org/10.6028/NIST.SP.800-82r3
13. K. Willcox and B. Segundo, "The role of computational science in digital twins," *Nature Comput. Sci.*, vol. 4, no. 3, pp. 147–149, 2024, doi: 10.1038/s43588-024-00609-4.
14. M. Intizar Ali et al., "Cognitive digital twins for smart manufacturing," *IEEE Intell. Syst.*, vol. 36, no. 2, pp. 96–100, Mar./Apr. 2021, doi: 10.1109/MIS.2021.3062437.

SABAH SUHAIL is a research fellow at the Centre for Secure Information Technologies (CSIT), Queen's University Belfast, Belfast, BT7 1NN, U.K. Her research interests include utilizing digital twins for cybersecurity and ensuring the security of digital twins within cyberphysical systems. Suhail received her Ph.D. degree in computer engineering from Kyung Hee University, South Korea. Contact her at s.suhail@qub.ac.uk.

MUBASHAR IQBAL is a lecturer of information security at the Institute of Computer Science, University of Tartu, 51009, Tartu, Estonia. His research interests include exploring security facets of blockchain systems; secure development of decentralized applications; and analyzing and modeling secure applications for the metaverse, digital twins, and intelligent infrastructures. Iqbal received his Ph.D. degree in computer science from the University of Tartu, Estonia. He is a Member of IEEE. Contact him at mubashar.iqbal@ut.ee.

KIERAN MCLAUGHLIN is a reader at the CSIT, Queen's University Belfast, Belfast, BT7 1NN, U.K., where he leads research in cybersecurity for operational technologies. His research interests include threat analysis, intrusion detection and prevention, and digital twins in cybersecurity contexts. McLaughlin received his Ph.D. degree in electronic engineering and computer science from Queen's University Belfast, U.K. Contact him at kieran.mclaughlin@qub.ac.uk.