

Systematization of Knowledge on Air-Gap Attacks and Defense Metrics: A Cyber Kill Chain Perspective

Bonn Park[†] , *Member, IEEE*, JiEun Lee , *Member, IEEE*, JiHo Lee , *Member, IEEE*,
Tarik Taleb , *Senior Member, IEEE*, and JaeSeung Song^{*} , *Senior Member, IEEE*

Abstract—This Systematization of Knowledge (SoK) consolidates research on air-gap attacks and defenses through a Cyber Kill Chain (CKC) perspective. We (i) classify physical-layer channels into three operational groups: electrical and electromagnetic (including radio-frequency emissions, optical signaling, and power-line conduction), mechanical (including airborne acoustic signals and structure-borne vibration), and thermal, (ii) analyze malware behaviors observed in isolated and industrial control system (ICS) environments, and (iii) align defense strategies with attack stages via an Air-Gap Cyber Kill Chain (ACKC) and its accompanying Defense Matrix.

Compared with prior surveys that emphasized single channels or ad-hoc demonstrations, our SoK models the entire attack process—from infiltration to trigger-based execution and exfiltration—highlighting operational prerequisites and potential defender intervention points. We clarify the notion of an “air-gap attack” relative to existing definitions and identify three core research gaps: fragmented cross-domain coverage, absence of quantitative detection metrics, and limited information technology (IT)–operational technology (OT) integration.

Building on these findings, we propose a unified, phase-wise mapping of offensive and defensive techniques applicable to closed networks. This SoK not only extends the CKC paradigm to physically isolated environments but also provides a metric-informed framework and a proposed reference evaluation configuration for future empirical studies. The contribution is analytical; this SoK does not report a new empirical defense benchmark or establish a standardized evaluation procedure.

Index Terms—Air-gap, CKC, C&C, Covert channel, ICS, TEMPEST, SoK.

I. INTRODUCTION

Air-gapped networks are physically isolated from external connectivity and are widely deployed in critical infrastructures such as defense, finance, and industrial control systems (ICS). Despite this isolation, recent advances have demonstrated that air-gapped systems can still be compromised through unconventional physical channels, which this Systematization of Knowledge (SoK) organizes into electrical and electromagnetic, mechanical, and thermal groups [1–5]. Stuxnet and Industroyer2 motivate analysis of industrial com-

promise; their inclusion does not establish that both incidents crossed a verified physical air gap [6, 7].

This observation reveals that existing air-gap studies lack a structured analytical framework capable of systematically capturing how attacks penetrate, persist, and operate across the entire life cycle of closed environments.

Frameworks such as the Cyber Kill Chain (CKC) [8] offer phase-based interpretations of cyberattacks, but most prior air-gap research has focused narrowly on demonstrating exfiltration channels rather than modeling full attack life cycles. As a result, existing literature lacks a unified perspective connecting infiltration vectors, malware behavior, command-and-control (C&C), and covert-channel operations.

To address this gap, we adopt a SoK approach that goes beyond a narrative survey by systematically extracting, comparing, and restructuring operational knowledge from prior air-gap attack studies and documented ICS malware incidents.

Building on these observations, we synthesize the Air-Gap Cyber Kill Chain (ACKC) as an analytical abstraction derived from recurring patterns identified across the literature and historical ICS attack cases, together with an associated defense matrix. This enables a clearer understanding of how attackers bypass isolation controls and how defenders can respond at each stage.

The main contributions of this SoK study are as follows:

- We reorganize fragmented air-gap attack research into a unified, process-oriented analytical model, providing a systematized view rather than introducing new attack mechanisms.
- We extract air-gap-relevant behaviors from documented ICS incidents and distinguish observed behavior from inferred applicability to physically isolated environments.
- We classify air-gap attack techniques by physical transmission media and analyze their operational constraints.
- We present a metric-informed defense matrix and a scenario-based validation protocol aligned with the ACKC, linking stage-specific observable indicators, defensive actions, validation scenarios, and evaluation metrics for closed networks.

This paper combines a broad survey of prior air-gap attack research with a SoK approach, reorganizing existing findings into a unified, life-cycle-based framework.

The remainder of this paper is organized as follows: Section II reviews the definitions of air-gapped systems, closed-network environments, and existing air-gap attack studies. Sec-

Bonn Park, JiEun Lee, JiHo Lee, and JaeSeung Song are with Sejong University, Seoul 05006, South Korea (e-mail: 23170993@sejong.ac.kr; je.lee@sejong.ac.kr; twozio@sejong.ac.kr; jssong@sejong.ac.kr).

Tarik Taleb is with the Faculty of Electrical Engineering and Information Technology, Ruhr University Bochum, 44801 Bochum, Germany (e-mail: tarik.taleb@rub.de).

[†]Bonn Park is the first author.

^{*}Prof. JaeSeung Song is the corresponding author.




Fig. 1. Schematic overview of previously studied air-gap attacks

tion III examines representative CKC and advanced persistent threat (APT) models to identify their structural assumptions and limitations. Section IV analyzes the characteristics of ICS malware relevant to air-gap attack stages. Based on these analyses, Section V derives the ACKC as a systematized life-cycle abstraction for air-gapped environments. Section VI categorizes physical-layer attack channels and presents the Air-Gap Defense Matrix aligned with the ACKC. Finally, Section VII summarizes the findings, discusses limitations, and outlines future research directions.

II. BACKGROUND KNOWLEDGE AND MOTIVATION

This section establishes the definitions, architectural assumptions, and analytical scope adopted throughout this SoK. Rather than introducing new motivations, it clarifies how prior studies operationalize air-gapped environments and which assumptions are implicitly embedded in existing evaluations. This foundation enables consistent interpretation of the literature and supports the limitation analysis that motivates the life-cycle-oriented modeling developed in subsequent sections. The evidence base is a focused, qualitative selection of the works cited in this manuscript, rather than an exhaustive systematic review. Sources are included for their relevance to physical-channel mechanisms, attack-life-cycle modeling, documented industrial or removable-media-mediated compromise, or corresponding defensive controls. Channel experiments, framework papers, and campaign reports serve different analytical roles and are not treated as interchangeable forms of validation.

The resulting synthesis is bounded by this cited evidence base. It does not support estimates of literature coverage or attack prevalence. A reproducible database search and screening record is not provided, which limits the reproducibility

of literature selection. The rationale for selecting the malware cases is described separately in Section IV.

A. Definition of Air-Gap

An air-gap refers to a security measure in which a closed network is physically isolated from external connectivity to protect sensitive systems from remote compromise. Such configurations are widely deployed in military, ICS, and financial environments to ensure data integrity and confidentiality. Despite their strong isolation guarantees, air-gapped systems remain vulnerable to indirect and multi-stage attacks that exploit physical proximity, compromised insiders, or covert physical channels.

An air-gap attack is therefore defined as any method that crosses the isolation boundary via unconventional or covert channels to enable information infiltration or data exfiltration [9]. Historically, concerns regarding unintentional information leakage date back to classified research such as *TEMPEST: A Signal Problem* by the U.S. National Security Agency (NSA) [10], which recognized that electromagnetic emanations could compromise confidentiality and laid the foundation for modern emission security (EMSEC).

B. Network Separation

This SoK distinguishes physical air gaps from other forms of network separation. A physically air-gapped environment has no wired or wireless network communication path across the specified isolation boundary, although communication among systems inside that boundary may remain available.

Logical isolation restricts communication through access controls or virtualization on shared infrastructure. Network

segmentation divides a network into domains whose inter-connections are controlled by policy. Virtual Local Area Networks (VLANs) can support logical separation, whereas Virtual Desktop Infrastructure (VDI) provides a virtualized desktop environment and does not, by itself, establish a physical air gap [11, 12].

For this paper, semi-isolated environments are deployments that permit limited or intermittent external connectivity, for example during authorized maintenance. These terms describe different operational conditions and are not necessarily mutually exclusive architectural categories.

The primary scope of the ACKC is physical network isolation. Cases involving logically isolated, segmented, or semi-isolated ICS environments are included to examine transferable behaviors, such as maintenance-mediated delivery, internal reconnaissance, and condition-dependent execution. Their inclusion does not imply that each case represents a verified breach of a physical air gap.

C. Existing Research on Air-Gap Security

Extensive academic research, including several recent studies, has demonstrated that air-gapped systems are vulnerable to information leakage through unconventional physical channels. This SoK organizes the channels examined in these studies into three operational groups: electrical and electromagnetic, mechanical, and thermal. The grouping reflects the principal mechanism used to convey and recover information [1–5, 13–16].

Across both earlier and more recent work, most studies focus on demonstrating the technical feasibility of crossing isolation boundaries under controlled experimental conditions, typically operating under the assumption that malicious code is already present within the isolated environment. Recent advances further show that software-only mechanisms, such as central processing unit (CPU) workload modulation or memory-access manipulation, can generate detectable electromagnetic emissions without relying on dedicated peripheral devices [14, 15].

Collectively, these studies indicate that air-gap attacks are inherently multi-modal in practice, with different physical media enabling information leakage under varying environmental and operational constraints. Figure 1 provides a high-level schematic overview of previously studied air-gap attack vectors, abstracting their common structural assumptions while highlighting the diversity of physical propagation mechanisms explored in the literature. The electrical and electromagnetic group includes radio-frequency (RF) emissions, optical signaling, and power-line conduction. Optical signaling includes visible-light and infrared signals. Power-line conduction conveys information through electrical variations on shared power wiring and is distinguished from radiated RF and optical emissions [1, 2, 5].

The mechanical group includes airborne acoustic signals and structure-borne vibration. These channels are distinguished by their propagation paths and receiving sensors: acoustic channels use sound transmitted through air, whereas vibrational channels use mechanical motion transmitted through structures [3, 13].

The thermal group includes temperature-modulation channels. Heat can be transferred by conduction through matter, convection involving fluid motion, and thermal radiation. Radiation is electromagnetic physically; here a channel is classified by its information-bearing variable and receiving measurement. Direct optical-intensity modulation is treated as optical signaling, whereas a receiver decoding temperature changes is treated as thermal signaling. The dominant heat-transfer mechanism depends on the specific deployment; this grouping does not establish that every reported thermal channel uses all three mechanisms.

This hierarchy is an operational grouping for channel analysis and defense mapping, rather than a mutually exclusive classification of all underlying physical phenomena. Electrical conduction along power wiring is distinct from thermal conduction. Inclusion of a power-line channel assumes a shared conductive path across the network-isolation boundary; it does not imply that every air-gapped deployment provides such a path. Figure 11 summarizes the three groups and their constituent channels.

We note that the experimental literature on physical covert-channel demonstrations cited in this section is concentrated within a small number of research groups, most notably Guri et al., who have produced the majority of systematically documented channel demonstrations across the electrical/electromagnetic, mechanical, and thermal groups. This concentration reflects the composition of the channel-demonstration sub-literature rather than the composition of this SoK as a whole: the malware-behavior analysis (Section IV), kill-chain and threat-intelligence framework literature (Section III), and reported ICS incidents (Section IV) draw on a broader set of academic and industrial sources, including MITRE, the NSA, ESET, Kaspersky, Mandiant, Presekal et al., Kour et al., and Kim et al. We removed redundant citations to the same group’s work where a single already-cited source was sufficient to support a given claim.

D. Limitations of Prior Studies and Open Challenges

The physical-channel studies reviewed in this SoK provide evidence of covert communication through different physical mechanisms, including RF emissions, optical signaling, thermal modulation, and power-line conduction [1, 2, 4, 5]. Their experimental results establish channel feasibility under the reported hardware and environmental conditions. Such results should be distinguished from evidence that a defense can reliably detect or suppress these channels under different deployment conditions.

The ESET, Kaspersky, and Mandiant reports used in the malware analysis provide complementary evidence about documented campaign behavior, including file collection, removable-media use, and offline or intermittent transfer opportunities [17–19]. These reports support the extraction of behaviors relevant to air-gap modeling. Their inclusion does not establish that every campaign crossed a verified physical air gap or that the proposed defenses have been empirically validated.

Despite their technical contributions, most prior studies on air-gap security share several structural limitations. First, they

commonly assume that malware has already been successfully deployed within the air-gapped environment, while infection vectors such as removable media, insider actions, or supply-chain compromise are treated only implicitly [14, 15]. This assumption simplifies experimental design but limits the ability to reason about early-stage attacker behavior and preventive defense mechanisms.

Second, many evaluations are conducted under idealized laboratory conditions that abstract away operational constraints present in real-world deployments, such as background electromagnetic noise, physical shielding, heterogeneous hardware configurations, and active monitoring systems [16]. As a result, detection thresholds and signal reliability reported in controlled settings may not directly translate to operational environments.

Moreover, recent advances in electromagnetic-based air-gap attacks, including software-only signal generation and frequency-hopping emission techniques, primarily concentrate on the signal generation and data exfiltration stages [14–16]. Earlier phases such as initial infiltration, long-term persistence, and coordinated multi-stage attack planning remain comparatively underexplored.

As illustrated in Fig. 1, existing research often models air-gap attacks as compact sequences centered on signal generation and data exfiltration. While such abstractions are effective for demonstrating technical feasibility, they frequently omit layered defenses, defensive feedback loops, and attacker decision-making across multiple stages.

These limitations indicate that air-gap threats should be analyzed as end-to-end, multi-phase processes encompassing infiltration, activation, alternatives to persistent command and control (C&C), and exfiltration under realistic constraints. Accordingly, the subsequent sections introduce a life-cycle-oriented analytical abstraction—the ACKC—that makes these assumptions explicit and supports phase-wise reasoning about defensive coverage in closed-network environments.

III. RELATED WORK AND COMPARATIVE POSITIONING

This section positions the proposed SoK against representative air-gap studies, covert-channel demonstrations, CKC-based attack models, APT life-cycle frameworks, and threat modeling studies for ICS and Cyber-Physical Systems (CPS). Rather than providing a tutorial-style review of individual frameworks, we examine their structural assumptions, operational scopes, and defense dependencies in order to clarify why existing models remain insufficient for physically isolated or air-gapped environments.

Before deriving an air-gap-specific kill chain, it is necessary to distinguish two research streams that have largely evolved separately. Prior air-gap and covert-channel studies provide detailed insights into physical-layer exfiltration channels, but they often assume that malware is already present inside the isolated system. Conversely, CKC and APT life-cycle models provide useful abstractions for staged cyber intrusions, but they typically rely on network connectivity, persistent command channels, and traffic visibility.



Fig. 2. Phases of the F2T2EA process




Fig. 3. Lockheed Martin CKC

A. Origins of the Kill Chain Concept

The concept of the *kill chain* originates from military doctrine, particularly within the U.S. Air Force’s time-sensitive targeting (TST) operations. This abstraction was designed to support decision-making against high-value or mobile targets by decomposing engagements into sequential phases. As illustrated in Fig. 2, this process was formalized as the Find, Fix, Track, Target, Engage, and Assess (F2T2EA) cycle, which emphasizes continuous situational awareness, real-time sensing, and centralized command authority [20, 21].

While effective in kinetic and network-enabled contexts, the F2T2EA model implicitly relies on persistent observability and timely command execution—assumptions that do not hold in physically isolated or air-gapped cyber environments.

B. Cyber Kill Chain and Network-Centric Assumptions

Building upon its military foundations, the kill chain concept was adapted to the cyber domain to address increasingly sophisticated, multi-stage intrusions. Lockheed Martin introduced the CKC, which decomposes cyber attacks into seven sequential phases, as shown in Fig. 3: reconnaissance, weaponization, delivery, exploitation, installation, C&C, and actions on objectives [22]. A key insight of this model is that disrupting any single phase may prevent the attack from progressing.

Subsequent studies refined the CKC abstraction. Verdasys emphasized lateral movement and argued that effective defense requires coordinated mitigation across multiple phases [23].

TABLE I
EXTENDED COURSE-OF-ACTION MATRIX FOR CKC (ADAPTED FROM PRIOR CKC-BASED DEFENSE LITERATURE)

Phase	Detection	Deny	Disrupt	Degrade	Deceive
Reconnaissance	Web analysis	FW ACLs	Timeout rules	Bandwidth throttle	Fake DNS, decoy host
Weaponization	NIDS	NIPS	Sandbox	Delay delivery	Fake payload repository
Delivery	Phishing alert, mail logs	Proxy filter	Inline AV	Queuing	URL morphing
Exploitation	HIDS, syscall anomaly	Patch management	DEP, code signing	Virtual execution	Dummy vulnerable service
Installation	HIDS	Whitelisting, chroot jail	AV/EDR	Rate limiting	Fake installer
C&C	NIDS	Egress blocking	NIPS, sinkholing	Tarpit, latency injection	DNS redirection
Actions on Objectives	Audit logs, DLP	File lockdown	USB blocking	QoS limiting	Honeypot, decoy database

Hewlett Packard Enterprise (HPE) further decomposed the final stage into endpoint-level sub-phases, such as staging and exfiltration, to provide finer-grained visibility into post-compromise behavior [24].

These models provide useful abstractions for intrusion analysis, but their application to isolated systems requires explicit treatment of connectivity, local execution, and available defender observations. These deployment conditions should not be interpreted as universal prerequisites of the CKC abstraction.

C. Course-of-Action Matrix as a Comparative Baseline

To operationalize the CKC for defense planning, Lockheed Martin proposed a Course-of-Action (CoA) matrix based on the D6 doctrine: Detection, Denial, Disruption, Degradation, Deception, and Destruction [25]. In enterprise contexts, the destruction phase is typically excluded, resulting in a five-dimensional defensive abstraction.

To ground our analysis in established CKC-based defense practice, we summarize a representative CoA matrix derived from prior work. Table I is intentionally retained as an analytical baseline rather than as a prescriptive defense framework. It illustrates how conventional cyber defenses have been historically structured around network-centric observability, centralized monitoring, and protocol-level inspection [22–24].

For clarity, we briefly expand the abbreviations used in Table I at their first occurrence. These terms correspond to commonly deployed, network-centric security controls assumed in CKC-based defense models.

Specifically, firewall (FW) and access control list (ACL) mechanisms are used for traffic filtering and access enforcement. Network intrusion detection system (NIDS) and network intrusion prevention system (NIPS) monitor network traffic for malicious activity, while host intrusion detection system (HIDS) operates at the endpoint level. Additional controls include antivirus (AV) software, endpoint detection and response (EDR), data execution prevention (DEP), and data loss prevention (DLP). Domain Name System (DNS) and quality of service (QoS) are also referenced as representative network-layer mechanisms for redirection, deception, or traffic regulation.

In air-gapped environments, persistent C&C channels and real-time traffic visibility are absent or severely constrained. Accordingly, the CoA matrix serves as a reference point to highlight the structural mismatch between traditional network-dependent defense assumptions and the operational realities of physically isolated systems.

D. ICS, Endpoint, and APT-Oriented Attack Frameworks

The MITRE ATT&CK for ICS framework provides a comprehensive taxonomy of adversarial tactics and techniques tailored to industrial environments [26]. Unlike CKC-based models, ATT&CK emphasizes behavioral enumeration rather than temporal attack progression. While valuable for threat intelligence sharing and post-incident analysis, the framework lacks an explicit life-cycle abstraction, and many techniques implicitly assume network-based access or remote interaction.

Endpoint-centric and APT-oriented life-cycle models address some of these limitations by providing finer-grained descriptions of post-compromise behavior. Hewlett Packard Enterprise’s Attack Life Cycle (ALC) model extends CKC concepts by emphasizing endpoint-level processes, including persistence, staging, and lateral movement [24]. Similarly, the APT life-cycle model proposed by Hassan II University emphasizes stealth, camouflage, and trace removal across multiple phases [27]. These characteristics align with certain air-gap malware traits, such as delayed activation and autonomous execution.

These frameworks provide complementary technique and life-cycle views. ATT&CK for ICS already includes removable-media entry under T0847. The analytical task here is to connect such behaviors to isolation conditions, physical-channel mechanisms, and phase-specific defensive opportunities [26].

E. Defense-Oriented CKC Extensions

Recent studies have extended CKC-based analysis beyond descriptive attack-stage modeling toward defense-oriented evaluation. Kour et al. integrated game theory with the CKC to model attacker–defender strategy selection across attack stages [28]. Their model represents attackers and defenders as

players in a non-cooperative mixed-strategy game, where attackers select stage-specific offensive strategies, while defenders select actions such as monitoring, detection, and response. This approach demonstrates that CKC-based reasoning can support quantitative defense evaluation using expected payoff and strategy probabilities.

Moving Target Defense (MTD) is a proactive defense strategy that dynamically changes system configurations, attack surfaces, or network resources to increase attacker uncertainty and reduce exploitation opportunities. Kim et al. proposed MTD in Depth (MTDID), a multi-phased MTD framework aligned with the CKC and MITRE ATT&CK frameworks [29]. Their approach categorizes MTD sub-techniques, such as virtual internet protocol (IP) address shuffling, software diversity, and redundancy, according to digital artifacts and CKC phases. MTDID further evaluates both security benefits and performance degradation in an experimental testbed based on software-defined networking (SDN).

These defense-oriented extensions show that CKC-based models can support phase-wise and measurable defense analysis. However, they remain grounded in networked industrial or SDN environments and do not explicitly address air-gap-specific constraints, such as removable-media-mediated malware injection, autonomous execution without persistent C&C, or covert physical-channel exfiltration.

F. Cyber-Physical Power System Kill Chain

Presekal et al. proposed the Advanced Cyber-Physical Power System (ACPPS) Kill Chain to model APT progression in power-grid environments [30]. Their framework extends conventional CKC and ICS-oriented models by incorporating technology–operational technology (IT–OT) compromise, unauthorized control of OT systems, cascading failures, blackout, and recovery.

This domain-specific model demonstrates that kill-chain frameworks must reflect the operational characteristics of the target environment. However, the ACPPS Kill Chain focuses on power-system impacts rather than air-gap-specific constraints such as removable-media-mediated malware injection, autonomous execution without persistent C&C, and covert physical-channel exfiltration.

G. Comparative Synthesis of Prior Frameworks

To synthesize the above discussion, Table II compares representative frameworks and study groups with respect to the proposed SoK. The comparison focuses on four dimensions that are central to air-gap attack modeling: life-cycle coverage, air-gap channel taxonomy, defense analysis, and IT–OT integration. These dimensions are selected because air-gap attacks require not only staged attack progression, but also explicit modeling of physical transmission media, measurable defensive responses, and operational constraints across IT and OT environments.

Ratings are qualitative author assessments of the reviewed works. For life-cycle coverage, High denotes coverage across major attack stages or tactics, Medium denotes coverage of a subset, and Low denotes a focus on an individual mechanism.

For channel taxonomy, High denotes a structured treatment of multiple physical channels, Medium a limited classification or discussion, and Low little or no channel classification.

For defense analysis, High denotes explicit comparative or quantitative evaluation within the study’s stated setting; Medium denotes organized controls or phase-wise defense mapping without such evaluation; and Low denotes only passing or no defense discussion. For IT–OT integration, High denotes explicit analysis linking IT compromise to industrial processes or effects, Medium denotes consideration of both domains without a developed linkage, and Low denotes little or no cross-domain treatment. These ratings describe scope and analytical depth, not independently validated superiority.

As shown in Table II, existing CKC and APT life-cycle models provide useful staged abstractions. This SoK builds on these abstractions by examining offline entry, control availability, condition-dependent execution, and physical-channel requirements together with phase-specific defenses. Air-gap covert-channel studies provide richer channel taxonomies, but their analytical focus is often concentrated on signal generation and exfiltration feasibility after compromise. Recent CKC-based defense studies, including game-theoretic and moving-target-defense-oriented models, further demonstrate the value of quantitative and phase-wise defense evaluation; however, they are still grounded in connected or software-defined network environments. Similarly, CPS-oriented kill-chain models highlight the need for domain-specific attack modeling, but they focus on power-system impacts rather than air-gap-specific operational constraints. This comparison motivates the proposed ACKC as an air-gap-oriented life-cycle abstraction that integrates physical infiltration, stealthy data collection, autonomous execution, covert channels, and phase-wise defense mapping.

IV. MALWARE CORPUS AND TRAIT EXTRACTION FOR AIR-GAP MODELING

The ICS malware cases examined in this section are included for their behavioral relevance to air-gap modeling, rather than because every incident is known to have crossed a physical air gap. As distinguished in Section II-B, the documented environments may involve different forms of isolation. These cases provide insight into transferable behaviors, including malware delivery, internal reconnaissance, and condition-dependent execution.

The four derivation cases were chosen to contrast industrial sabotage, asset reconnaissance, credential-mediated compromise, and protocol-specific execution. The three additional campaigns emphasize removable-media transfer and offline collection, and constitute an illustrative applicability check rather than a blind or statistically independent validation set. Campaign membership does not establish a physical air-gap breach. Hardware implants, firmware-only compromise, and laboratory optical channels are outside this campaign-level check; physical-channel studies are examined separately in the channel analysis. Consequently, the corpus cannot validate coverage of all air-gap vectors or estimate their prevalence.

TABLE II
COMPARISON OF PRIOR AIR-GAP AND KILL-CHAIN STUDIES WITH THIS SoK

Work / Framework	Main Focus	Life-cycle Coverage	Air-Gap Channel Taxonomy	Defense Analysis	IT-OT Integration	Limitations for Air-Gap Modeling
Lockheed Martin CKC [22]	Enterprise APT life-cycle	High	Low	Medium	Low	Requires deployment-specific detail for isolation conditions and physical-channel analysis
MITRE ATT&CK for ICS [26]	ICS adversarial techniques	High	Low	Medium	High	Provides comprehensive technique coverage, but does not model air-gap-specific attack progression
HPE Attack Life Cycle [24]	Endpoint-level attack behavior	High	Low	Medium	Medium	Limited modeling of physical infiltration and covert exfiltration
APT Life Cycle Models [27]	Stealthy long-term intrusion	High	Low	Medium	Medium	Assumes networked attacker control and does not model air-gap channels
Air-Gap Covert-Channel Studies [31]	Physical-channel exfiltration	Medium	High	Medium	Low	Often assumes malware is already installed inside the isolated system
ICS Malware Case Studies [32]	Closed-network malware behavior	Medium	Low	Medium	High	Useful for trait extraction but insufficient for generalized ACKC validation
Game-Theoretic CKC Models [28]	Attacker-defender strategy evaluation	High	Low	High	Medium	Quantifies CKC-based defense strategies but does not model air-gap physical channels
MTD-Based CKC Defense Models [29]	Phase-wise proactive defense	High	Low	High	Medium	Evaluates CKC-aligned MTD techniques, but assumes networked SDN environments
Cyber-Physical Power System Kill Chain [30]	APT modeling for power-grid environments	High	Low	Medium	High	Captures IT-OT and physical impacts, but not air-gap covert exfiltration
This SoK	ACKC-based air-gap attack modeling	High	High	Medium	Medium	Integrates physical infiltration, autonomous execution, covert channels, and phase-wise defense mapping; defense analysis remains analytical pending empirical validation




Fig. 4. Classification of ICS Malware [32]

A. Derivation Corpus: Representative ICS Malware Cases

Human-mediated transfer is modeled across phases. An unwitting user may carry a contaminated device, a manipulated maintainer may bypass an intake procedure, and a malicious insider may exploit authorized access deliberately. Proposed controls combine role-appropriate personnel vetting, recurring social-engineering exercises using maintenance scenarios, least privilege, two-person approval for exceptional transfers, and reconciliation of user, media, and work-order logs. Behavioral alerts should concern deviations from approved tasks rather than treating every authorized access as hostile. These controls reduce exposure but do not eliminate trusted-insider risk.

Previous research categorizes ICS malware according to targeting characteristics, as shown in Fig. 4. Targeted malware is subdivided into *ICS-themed* and *ICS-tailored* families:

the former engages in reconnaissance or data manipulation without deep hardware specificity, whereas the latter is developed with explicit knowledge of industrial devices or communication protocols. A notable example is Stuxnet, which precisely targeted Siemens Step7-based programmable logic controllers (PLCs) to damage uranium centrifuges in Iran’s Natanz facility [6]. In contrast, Havex primarily focused on reconnaissance, while BlackEnergy3 and Industroyer2 demonstrated the capability for large-scale operational disruption. By comparing these representative cases, this section derives generalizable technical traits that underpin the proposed ACKC framework.

1) *Stuxnet Malware*: As illustrated in Fig. 5, Stuxnet, discovered in 2010, was a sophisticated ICS-targeted malware engineered to sabotage uranium enrichment centrifuges at Iran’s Natanz nuclear facility. It infiltrated air-gapped systems via infected Universal Serial Bus (USB) drives exploiting the Windows Link (LNK) vulnerability (CVE-2010-2568) for initial execution. After establishing persistence through privilege escalation, it propagated internally via the Print Spooler vulnerability (CVE-2010-2729). Once embedded within engineering workstations, it targeted Siemens Step7 PLCs, intercepting control commands and falsifying operator feedback to conceal physical damage. All referenced vulnerabilities follow the Common Vulnerabilities and Exposures (CVE) framework [33].

2) *Havex Malware*: As shown in Fig. 6, Havex (also known as Oldrea), first reported in 2013, primarily targeted ICS environments through supply-chain compromise and spear-phishing [34]. The campaign distributed a remote




Fig. 5. Stuxnet Attack Flow and Derived Characteristics of Air-Gap Malware




Fig. 6. Havex Attack Flow and Derived Characteristics of Air-Gap Malware

access trojan (RAT) that established HTTP(S)-based callbacks to its C&C servers. The component `mbcheck.dll` performed Object Linking and Embedding (OLE) for Process Control (OPC) enumeration and related ICS reconnaissance, enabling attackers to map devices and collect process/state data. OPC, an industry-standard communication interface, allows PLCs, sensors, and workstations to exchange real-time process data; Havex abused this legitimate protocol to locate and profile assets even within segmented networks. The full attack flow—initial infection, RAT deployment, C&C establishment, OPC enumeration, data exfiltration, and stealth persistence—aligns with the sequence depicted in the figure [35, 36].

3) *BlackEnergy3 Malware*: As illustrated in Fig. 7, the 2015 Ukraine blackout was attributed to the BlackEnergy3 campaign [37]. This modular malware was delivered via spear-phishing emails containing malicious Microsoft Office documents (e.g., `black.xls`), which deployed a dropper into the `%LOCALAPPDATA%` directory. Once executed, the dropper established HTTP(S) connections to its C&C servers to fetch additional modules and propagate laterally using stolen credentials. After compromising IT networks, it expanded to HMIs and engineering workstations bridging into OT domains, issuing unauthorized breaker commands and deploying the KillDisk wiper to erase critical files. This attack chain demonstrates three characteristics relevant to air-gap threat modeling: stealthy operation through process masquerading and C&C communication, payload execution from user-level directories, and ICS disruption through HMI/EWS access, unauthorized breaker commands, and KillDisk-based file wip-




Fig. 7. BlackEnergy3 Attack Flow and Derived Characteristics of Air-Gap Malware




Fig. 8. Industroyer2 Attack Flow and Derived Characteristics of Air-Gap Malware

ing [37].

4) *Industroyer2 Malware*: As shown in Fig. 8, Industroyer2, first reported in 2022, is a specialized malware family designed to directly manipulate electric substations by abusing the International Electrotechnical Commission (IEC) 60870-5-104 (IEC-104) protocol [38]. The malware contained hardcoded configurations and device mappings, allowing autonomous operation without persistent external C&C.

Industroyer2 was reported in 2022 in an attempted attack on Ukrainian electrical substations. Its hardcoded configuration and scheduled execution support analysis of condition-dependent industrial operations without continuous external control. The ESET report identifies the executable as `108_100.exe` and describes use of the International Electrotechnical Commission (IEC) 60870-5-104 protocol. It reports coordinated deployment of CaddyWiper, including scheduled wiping on the Industroyer2 host, but explicitly states that initial compromise and movement from the IT network into the ICS network were unknown. These unknown steps are not treated as established phishing, credential theft, or a verified physical air-gap crossing [38].

B. Extended Corpus: Recent Air-Gap-Targeting Campaigns

The preceding cases provide a derivation corpus for extracting recurring ICS and closed-network malware traits. However, to address more recent removable-media and offline-exfiltration threats, we further examine recent air-gap-targeting campaigns, including Ramsay, USBCulprit, and PlugX/SOGU USB propagation variants. This extended corpus is not used

as the sole basis for deriving the ACKC; rather, it is used to examine whether the extracted traits remain observable in more recent air-gap-targeting campaigns.

1) *Ramsay*: Ramsay is a cyber-espionage toolkit reported in 2020 and specifically analyzed as being tailored for air-gapped networks [17]. Unlike destructive ICS malware, Ramsay primarily emphasizes stealthy document discovery, collection, and staging inside isolated environments. Its operational logic is relevant to air-gap modeling because it does not rely on continuous network connectivity after compromise; instead, it prepares collected artifacts for later retrieval through indirect or offline mechanisms.

From the ACKC perspective, Ramsay is most closely associated with the **Stealth and Data Collection** phase and the data-oriented branch of the **Attacks** phase. Its emphasis on sensitive file discovery and internal staging supports the observation that air-gap attacks may remain dormant for extended periods before exfiltration becomes possible. Therefore, Ramsay strengthens the need to model air-gap malware not only as a destructive payload, but also as an autonomous collection framework operating under constrained connectivity.

2) *USBCulprit*: USBCulprit is a USB-borne air-gap malware case associated with the Cycldek activity set and analyzed in both security reporting and academic literature [18, 39]. It is particularly relevant to air-gap attack modeling because it relies on removable media as both an infection and data-transfer mechanism. This behavior directly reflects the operational role of USB devices as a bridge between connected and isolated environments.

In ACKC terms, USBCulprit maps strongly to the **Malware Injection** phase because the removable device functions as the delivery vector into otherwise isolated systems. It also contributes to the **Stealth and Data Collection** phase by harvesting files from compromised hosts and preparing them for later transfer. The case illustrates why air-gap defense must include media-control policies, device identity tracking, quarantine procedures, and scan-latency metrics rather than relying only on network monitoring.

3) *PlugX/SOGU USB Propagation Variants*: Recent PlugX/SOGU USB propagation variants demonstrate that removable-media-based attack paths remain operationally relevant even in modern environments [19, 40]. These campaigns use infected USB drives to propagate between hosts and support espionage-oriented collection, showing that physical transfer paths continue to complement network-based compromise. Unlike purely network-centric malware, USB-propagating variants can move across disconnected or intermittently connected systems through normal user and maintenance workflows.

From the ACKC perspective, PlugX/SOGU USB variants reinforce three recurring air-gap traits: removable-media-mediated injection, autonomous propagation, and delayed exfiltration. They primarily map to the **Malware Injection** and **Stealth and Data Collection** phases, while their offline staging behavior is also relevant to the data-exfiltration branch of the **Attacks** phase. This extended corpus confirms that the ACKC is not limited to early ICS malware such as Stuxnet or

Havex, but remains applicable to more recent USB-mediated air-gap-targeting campaigns.

Table III summarizes how recent air-gap-targeting campaigns instantiate recurring behavioral traits observed in the derivation corpus. These cases demonstrate that removable-media-mediated injection, autonomous operation, stealthy data collection, and delayed exfiltration remain relevant in modern air-gap attack scenarios.

C. Comparative Summary of Air-Gap Characteristics

The analyzed corpus shows that air-gap-capable malware is better characterized by recurring operational traits than by taxonomy alone. Prior ICS malware categorization distinguishes non-targeted, ICS-themed, and ICS-tailored malware [32]; however, air-gap attack modeling requires identifying behavioral traits that remain observable across both representative ICS malware and recent air-gap-targeting campaigns.

The derivation corpus, consisting of Stuxnet, Havex, Black-Energy3, and Industroyer2, provides representative examples of malware behavior in closed or semi-isolated ICS environments [6, 34, 37, 38]. These cases show recurring traits such as physical or offline malware injection, autonomous operation under constrained connectivity, stealthy internal reconnaissance, staged data collection, and abnormal OT or hardware interaction. The extended corpus, including Ramsay, USBCulprit, and PlugX/SOGU USB propagation variants, further demonstrates that these traits remain relevant in more recent air-gap-targeting campaigns, especially those involving removable media, offline collection, artifact staging, and delayed transfer [17–19, 39, 40].

As summarized in Table III, the extended corpus instantiates the behavioral traits extracted from the derivation corpus in more recent USB-mediated and offline-exfiltration scenarios. By comparing both corpora, this section identifies five recurring traits relevant to ACKC modeling: (i) physical or offline malware injection, (ii) autonomous operation, (iii) stealthy data collection, (iv) offline or covert exfiltration, and (v) abnormal OT or hardware interaction.

Importantly, trigger or condition-aware logic is not treated as a standalone trait mapped exclusively to a single phase. Instead, such logic may be implemented during malware creation, remain dormant during stealth and data collection, and manifest operationally during the attacks phase when pre-defined time-, event-, or system-state conditions are satisfied. This distinction separates where the logic is designed, where it remains latent, and where its operational effect becomes observable. Malware Creation denotes attacker-side preparation. Defender inspection of trigger logic becomes possible only when an artifact is available for intake analysis or after compromise; it is retrospective evidence about preparation, not direct observation of the attacker developing malware.

These recurring traits collectively form the analytical foundation for designing the ACKC presented in Section V, offering a structured understanding of how air-gap attacks evolve across multiple stages and how corresponding defenses can be systematically derived.

TABLE III
RECENT AIR-GAP-TARGETING CAMPAIGNS AND EXTRACTED ACKC TRAITS

Campaign	Targeting Context	Air-Gap-Relevant Attack Path	Extracted Behavioral Trait	ACKC Phase Mapping
Ramsay [17]	Air-gapped cyber-espionage environment	Offline operation after compromise and artifact staging for later retrieval	Sensitive-document discovery, internal collection, and delayed exfiltration preparation	Stealth and Data Collection; Attacks
USBCulprit [18, 39]	Isolated or semi-isolated hosts reachable through removable media	USB-mediated infection, file harvesting, and physical transfer across isolation boundaries	Removable-media-mediated injection and USB-assisted data movement	Malware Injection; Stealth and Data Collection
PlugX/SOGU USB variants [19, 40]	USB-enabled espionage and host-to-host propagation environments	Infected USB drives, autonomous propagation, offline staging, and delayed transfer through user or maintenance workflows	Autonomous propagation, offline staging, and delayed exfiltration under intermittent connectivity	Malware Injection; Stealth and Data Collection; Attacks

V. TOWARDS AIR-GAP CYBER KILL CHAIN

This section derives the ACKC as a SoK-based abstraction rather than a new operational doctrine. Building on the comparative review in Section III and the malware traits extracted in Section IV, we organize the extracted behaviors into an air-gap-oriented analytical view and examine their implications for defensive intervention. Unlike conventional IT intrusions, air-gap operations must explicitly account for (i) physical or offline malware injection, (ii) constrained or absent command and control (C&C), and (iii) signal-mediated execution and exfiltration [14–16, 31, 41].

The Lockheed Martin CKC provides a sequential intrusion abstraction [22], whereas MITRE ATT&CK for ICS offers a fine-grained taxonomy of adversarial tactics and techniques for industrial environments [26]. HPE’s ALC complements these perspectives with an endpoint-centric focus on post-compromise persistence and lateral behaviors [24]. The APT Life Cycle model proposed by Hassan II University highlights camouflage and trace removal as persistent properties across stages [27]. When applying these models to an isolated deployment, the analyst must specify which communication paths and observation sources are actually available. Accordingly, an air-gap-oriented kill chain should preserve the explanatory advantages of phase-based reasoning while explicitly incorporating offline propagation and covert-channel operations.

The ACKC organizes physical infiltration, trigger-based execution, and subsequent operations under isolation constraints within a common analytical structure. These behaviors are not claimed to be unrepresentable in the original CKC. Compared with MITRE ATT&CK for ICS, it emphasizes end-to-end progression (from preparation to execution) rather than enumerating isolated techniques. Compared with HPE’s ALC and the APT Life Cycle, it reinterprets persistence and camouflage under non-continuous control where external interaction may be intermittent, indirect, or receiver-mediated.

Ultimately, the ACKC aims to support proactive mitigation by clarifying when and where defenders can intervene before attacker objectives materialize in a closed network. The following subsections present the objectives of adopting a kill-chain lens, analyze limitations of prior models for isolated systems, and introduce the proposed ACKC.

A. Objectives of Air-Gap Cyber Kill Chain

Air-gap attacks are challenging because they combine sophisticated infiltration with covert or unconventional transmis-

sion methods. Nevertheless, their potential impact on critical systems motivates structured, phase-based threat modeling. We summarize four objectives for adopting a CKC perspective in air-gapped environments.

(1) **End-to-end interpretation of attack progression.** A kill chain perspective enables systematic interpretation from external preparation to internal operation and final execution, rather than assuming malware presence inside a closed network as a starting point.

(2) **Earlier detection and proactive response.** By explicitly representing pre-injection and pre-execution phases, defenders can monitor preparatory indicators and incorporate extracted traits into detection pipelines.

(3) **Phase-wise exposure and vulnerability assessment.** Decomposing the life cycle supports identification of weak links across people, process, and technology controls, including media handling, engineering workstation hardening, and segmentation boundaries between IT and OT.

(4) **Post-incident reconstruction and resilience.** After compromise, phase-based reasoning supports forensic reconstruction of which controls failed and why, enabling targeted reinforcement rather than ad-hoc hardening.

B. Analysis of Cyber Kill Chain Models for Isolated Systems

Building on the comparative review in Section III, we reinterpret prior kill-chain and attack-life-cycle models under isolation constraints. The original CKC includes removable-media delivery, and ATT&CK for ICS includes T0847, Replication Through Removable Media. Offline entry therefore does not, by itself, establish a new modeling capability of the ACKC. The comparison below examines how isolation conditions and physical-channel requirements are organized within the proposed phase-based analysis [22, 26].

The following mapping is many-to-many and analytical, rather than a claim that ACKC replaces CKC or ATT&CK for ICS. External preparation has no necessary one-to-one counterpart in the ICS technique matrix; internal discovery after entry is placed under Stealth and Data Collection.

- 1) Closed Network Reconnaissance corresponds to CKC Reconnaissance, with attention to maintenance schedules, suppliers, and access to the isolation boundary. T0846 (Remote System Discovery) and T0842 (Network Sniffing) describe internal discovery after access and therefore belong primarily to Stealth and Data Collection

in this mapping, not automatically to external reconnaissance.

- 2) **Malware Creation** corresponds to CKC Weaponization. Offline trigger logic and target-specific configuration are preparation properties; no direct ICS technique identifier is assigned to malware development itself.
- 3) **Malware Injection** corresponds to CKC Delivery and, where applicable, Exploitation and Installation. T0847 (Replication Through Removable Media) is a representative ICS technique, but does not cover every maintenance or insider transfer.
- 4) **Stealth and Data Collection** overlaps CKC Installation and preparatory Actions on Objectives. T0846, T0842, and T0811 (Data from Information Repositories) illustrate discovery and collection within the protected environment. Persistent external command and control is not required.
- 5) **Attacks** corresponds to CKC Actions on Objectives. T0809 (Data Destruction) and T0882 (Theft of Operational Information) illustrate different objectives. T0882 identifies information theft, not a particular optical, acoustic, or other physical transfer mechanism; that mechanism is recorded separately in ACKC.

The mapping demonstrates substantial overlap with existing frameworks. The contribution of the ACKC is the systematic integration of isolation-boundary crossing, control availability, local execution conditions, and physical-channel requirements with defensive opportunities. These details could also be documented through annotations or extensions to existing frameworks. The proposed synthesis does not establish greater expressive power or improved detection performance.

C. Proposed Air-Gap Cyber Kill Chain Model

Building on the preceding analysis, we propose the ACKC, which organizes air-gap attack progression into five sequential phases: *Closed Network Reconnaissance*, *Malware Creation*, *Malware Injection*, *Stealth and Data Collection*, and *Attacks*. The ACKC explicitly considers trigger-based, intermittent, and signal-mediated operations under physical isolation, while retaining correspondences with existing kill-chain stages [6, 14, 38, 39].

As illustrated in Fig. 9, the ACKC captures the full attack flow observed in closed industrial networks by integrating both physical infiltration and covert-channel execution paths. Each phase reflects operational constraints typical of air-gapped environments, where adversaries rely on indirect reconnaissance, removable media, and activation logic rather than continuous C&C channels.

The five phases of the ACKC are summarized as follows:

- **Closed Network Reconnaissance:** Identification of the target environment through open-source intelligence (OSINT), social engineering, or supply-chain analysis to collect configuration, workflow, and personnel information relevant to isolated systems.
- **Malware Creation:** Development of customized malware engineered for offline execution, environment-specific vulnerabilities, and protocol-/process-aware behavior within closed networks.

- **Malware Injection:** Physical or offline introduction of malware via removable media, engineering laptops, maintenance workflows, or insider-enabled transfer, which substitutes for network-based delivery.
- **Stealth and Data Collection:** Establishment of persistence, internal reconnaissance, and staged preparation that enables later execution while minimizing observable exposure. This phase may include asset discovery, credential harvesting, project file infection, and preparation of covert-channel parameters, but avoids overt exfiltration or destructive actions to reduce detection risk under constrained observability.
- **Attacks:** Execution of one or both optional objectives: system neutralization and data exfiltration. Neutralization may manipulate industrial logic or disrupt availability. Exfiltration may use delayed removable-media retrieval or covert channels in the electrical and electromagnetic, mechanical, or thermal group. The objectives may occur independently, in either order, or concurrently; neutralization is not a prerequisite for exfiltration. Activation may depend on preconfigured time, event, or system-state conditions, with no requirement for persistent external command and control.

VI. DEFENSE METHODS AGAINST AIR-GAP CYBER ATTACKS

Following the taxonomy established in Section II-C, this section organizes physical-channel analysis into electrical and electromagnetic, mechanical, and thermal groups. The detailed signal-based defense discussion focuses on RF emissions and acoustic signals as representative channels within the first two groups [1, 3, 15, 16, 42, 43].

The proposed defense matrix is intended to accommodate channel-specific observables and controls across all three groups, including optical signaling, power-line conduction, structure-borne vibration, and temperature-based signaling. These channels require different sensing and mitigation mechanisms according to their transfer paths and deployment conditions [2, 4, 5, 13]. Air-gapped systems provide a primary layer of security through physical isolation, yet this very isolation introduces unique vulnerabilities that conventional network-centric defenses cannot fully address. Prior research on air-gap attacks has illuminated various covert channels within the electrical and electromagnetic, mechanical, and thermal groups, but often presumes that air-gap malware is already installed, yielding a fragmentary view of the end-to-end attack life cycle [14, 15, 31]. To counter these threats effectively, defenders require a detailed understanding of both malware creation paths and sophisticated data exfiltration techniques.

As elaborated in Section V, this study adopts the ACKC as a structured lens for analysis. Building on this framework, we develop a defense paradigm along two axes: (i) methods derived from behavioral characteristics of air-gap malware, and (ii) mechanisms organized by physical transmission media. While the Lockheed Martin Course-of-Action Matrix is valuable for general cyber defense, it does not fully capture air-gap peculiarities such as autonomous operation and abnormal




Fig. 9. Overview of the proposed ACKC.

hardware control within closed networks. Therefore, we introduce a metric-informed Air-Gap Defense Matrix (Table IV) that aligns ACKC phases with observable indicators, defensive actions, and defense objectives.

Fig. 9 presents the ACKC as an attacker-side process flow. Fig. 10 overlays the corresponding defender-side view on the same five-step flow: each step is paired with a representative defensive action drawn from Table IV, so that the attacker-side progression in Fig. 9 and the defender-side intervention points in Fig. 10 can be read side by side. Table IV remains the authoritative, complete mapping of observable indicators, defensive actions, and defense objectives; Fig. 10 presents a subset of this mapping as a visual summary rather than a replacement for the table.

Because air-gapped environments differ in mission criticality, sensor availability, and acceptable operational disruption, the matrix does not prescribe universal numerical thresholds. Instead, Table V summarizes measurable outcome categories and operational cost factors that can be instantiated in site-specific evaluations. Accordingly, the proposed metrics should be interpreted as evaluation categories rather than universally applicable benchmark values. Their primary purpose is to provide a common analytical basis for comparing defense effectiveness across heterogeneous air-gapped environments while allowing deployment-specific calibration according to operational requirements.

A. Metric-Informed Air-Gap Defense Matrix

To operationalize the behavioral insights derived from prior air-gap malware analyses, we organize phase-specific defenses into the metric-informed Air-Gap Defense Matrix (Table IV). The matrix links each ACKC phase to observable indica-

tors, defensive actions, and defense objectives. The defensive actions are derived from five control dimensions: anomaly detection, media control, execution restriction, deception, and asset sanitization [22, 25]. For the Malware Creation phase, the listed observable indicators (e.g., unsigned binaries, hash mismatches) are necessarily detected only after the artifact enters the defender's environment, typically during the Malware Injection or Stealth and Data Collection phases. The phase label in Table IV reflects where the indicator originates in the attacker's process, not the point at which it becomes observable to defenders.

Table IV specifies phase-wise observable indicators and defensive actions, while Table V clarifies how these defenses can be evaluated through measurable outcomes. Here, mean time to detect (MTTD) denotes the average time between the occurrence of an attack behavior and its detection, while mean time to restore (MTTR) denotes the average time required to recover affected assets to a trusted operational state. For example, the media-control actions listed for the Malware Injection phase in Table IV can be evaluated using the corresponding metrics in Table V, such as unauthorized media block rate, removable-media scan latency, and quarantine success rate. Similarly, defenses in the Attacks phase can be evaluated using signal detection rate, false-alarm rate, leakage-bandwidth reduction, and restoration time. These metrics are not intended as universal thresholds; rather, they provide evaluation categories that can be calibrated according to mission requirements, sensor availability, and acceptable operational disruption. For instance, false-positive tolerance, acceptable scan latency, and recovery-time targets can be defined differently for mission-critical control systems, maintenance workstations, and monitoring-only assets.




Fig. 10. Application of the Air-Gap Defense Matrix to ACKC phases.

TABLE IV
ACKC-ALIGNED AIR-GAP DEFENSE MATRIX

ACKC Phase	Observable Indicators	Defensive Actions	Defense Objective
Closed Network Reconnaissance	Supplier impersonation; abnormal maintenance inquiries; exposed engineering assets; unusual access requests; OSINT-based targeting indicators	OSINT monitoring; supplier verification; access-pattern review; controlled disclosure of system information; honeypots and falsified configurations	Early detection of reconnaissance and deception of adversarial profiling
Malware Creation	Unsigned binaries; abnormal executable metadata; suspicious ICS libraries; hash mismatches; unauthorized scripts or configuration files	File-integrity monitoring; code-signing verification; offline sandbox analysis; application allowlisting; signed repositories; configuration-data deception	Reduction of custom malware feasibility and prevention of unauthorized code introduction
Malware Injection	New USB serial numbers; unauthorized removable-media mounts; file copies from maintenance laptops; autorun changes; execution from external drives	USB allowlisting; port lockdown; removable-media scanning; media quarantine; autorun disablement; bait media; device-access auditing	Prevention and containment of physical or insider-assisted malware delivery
Stealth & Data Collection	Abnormal file-access bursts; project-file enumeration; archive staging; hidden directory creation; unusual ICS query frequency; credential misuse	File-access anomaly detection; least-privilege enforcement; ICS command restriction; decoy project files; virtual decoy systems; persistence cleanup; log-integrity restoration	Detection and suppression of persistence, internal reconnaissance, and staged data collection
Attacks (Neutralization / Data Exfiltration)	Trigger-time activation; abnormal PLC commands; actuator state changes; RF/acoustic/optical emission peaks; periodic modulation patterns; abnormal firmware behavior	Logic restriction; interface control; execution isolation; emission shielding; spectrum monitoring; decoy emissions; controlled jamming; state restoration; system re-imaging	Mitigation of destructive execution, covert exfiltration, and restoration of trusted system state

B. Scenario-Based Validation Protocol

Although this SoK does not conduct a large-scale empirical benchmark, the proposed ACKC and Air-Gap Defense Matrix are designed to support reproducible validation in controlled air-gap testbeds. To make this evaluation path explicit, we outline a scenario-based validation protocol that maps representative air-gap attack behaviors to ACKC phases, observable data sources, defensive actions, and measurable outcomes. This protocol is intended to bridge the gap between conceptual kill-chain modeling and practical evaluation without requiring deployment of live malware in operational environments. Unlike the extended-corpus analysis in Section IV, which checks whether ACKC-relevant traits remain observable in

independent recent air-gap-targeting campaigns, the validation protocol in this subsection focuses on how ACKC-aligned defenses can be exercised and measured in controlled scenarios. It therefore complements Table IV, which defines phase-wise observable indicators and defensive actions, and Table V, which defines the corresponding evaluation metric categories.

The following is a proposed reference configuration for future evaluation, not a testbed implemented or measured in this SoK. Use four separately imaged computers: a transfer-preparation host outside the protected network, a disconnected intake/quarantine station, an engineering workstation, and a controller-simulator host. Use OpenPLC Runtime v4 on the last host with a software-only tank-level control program and

TABLE V
EVALUATION METRICS FOR ACKC-ALIGNED AIR-GAP DEFENSES

Metric Category	Example Metrics	Evaluation Meaning	Relevant ACKC Phase
Detection Effectiveness	Detection rate (recall); precision; false-positive rate; false-alarm rate	Measures how accurately suspicious behaviors, malicious files, abnormal commands, or covert signals are detected without excessive false alarms	All phases
Time Efficiency	MTTD; MTTR; analyst triage time; alert-to-containment time	Measures how quickly defenders detect, analyze, contain, and restore from phase-specific attack behaviors	All phases
Media Control	Unauthorized media block rate; removable-media scan latency; quarantine success rate; approved-media false-positive rate	Evaluates whether removable media and maintenance devices are effectively controlled without excessive disruption to legitimate operations	Malware Injection
Execution Restriction	Allowlist violation count; unauthorized execution block rate; code-signing violation count; engineering-tool false-positive rate	Measures the effectiveness and side effects of restricting unauthorized binaries, scripts, and execution paths	Malware Creation; Malware Injection; Stealth & Data Collection
Stealth Detection	File-access anomaly detection rate; hidden-directory detection rate; credential-misuse detection rate; abnormal ICS-query detection rate	Evaluates whether persistence, internal reconnaissance, staged collection, and covert preparation behaviors are detected	Stealth & Data Collection
Covert-Channel Suppression	Signal detection rate; leakage-bandwidth reduction; signal-to-noise ratio change; emission anomaly score; jamming effectiveness	Measures whether covert channels in the electrical and electromagnetic, mechanical, or thermal group are detected, degraded, or suppressed	Attacks (Data Exfiltration)
Recovery Effectiveness	System restoration time; re-imaging time; firmware integrity restoration rate; verified clean-state recovery rate	Evaluates how quickly and reliably compromised assets are restored to a trusted operational state	Attacks; post-incident recovery
Operational Cost	Analyst workload; maintenance delay; scanning overhead; sensor deployment cost; shielding cost; downtime	Captures the operational burden introduced by defense mechanisms and helps compare security benefit against deployment cost	All phases

simulated state variables; no physical actuators are connected. A fifth computer collects local logs and sensor recordings. The engineering workstation, simulator, and collector may communicate over an internal switch, but have no wired or wireless network path to the external host. Approved removable media provide the only administrative transfer route. Record operating-system and simulator versions, image hashes, hardware models, internal addressing, control period, and the allowed control operations in the run manifest. The selected controller platform is an implementation option, not evidence that this configuration has been tested. Within such an environment, each ACKC phase can be instantiated as a benign or emulated behavior corresponding to representative attack stages. For example, malware injection can be emulated through controlled removable-media file transfer, stealth and data collection can be represented by scripted file enumeration or archive staging, and the attacks phase can be represented by simulated PLC command anomalies or controlled signal emissions. Table VI operationalizes this protocol by mapping each representative validation scenario to an ACKC phase, observable data source, defense to evaluate, and primary metric. The resulting observations can then be evaluated using the metric categories in Table V, including detection rate, false-positive rate, MTTD, media-control latency, signal-detection rate, leakage-bandwidth reduction, recovery time, and restoration success rate.

As shown in Table VI, each validation scenario is tied to a specific ACKC phase and to data sources that can be collected without deploying live malware in operational environments. For instance, controlled removable-media transfer evaluates the Malware Injection phase through USB serial numbers, mount events, file-copy logs, and scan results, whereas scripted file enumeration and staging evaluates the Stealth and Data Collection phase through file-access bursts, archive creation, hidden directories, and credential-use logs. Similarly, simu-

lated ICS command anomalies and controlled signal-emission experiments instantiate the two branches of the Attacks phase, namely neutralization and data exfiltration. This protocol does not prescribe universal thresholds because acceptable false-positive rates, scan latency, recovery time, and operational disruption vary across mission-critical control systems, maintenance workstations, and monitoring-only assets. Instead, it provides a reproducible evaluation structure through which future studies can instantiate site-specific thresholds and compare defense coverage across ACKC phases. Accordingly, Table VI should be interpreted as a validation blueprint rather than an empirical benchmark table. It connects the ACKC phases, the observable indicators in Table IV, and the evaluation metrics in Table V, thereby providing a reproducible pathway for future controlled experiments and site-specific validation of air-gap defenses.

C. Behavioral Basis for ACKC-Aligned Defense

Conventional security mechanisms such as intrusion detection systems (IDSs), intrusion prevention systems (IPSs), and firewalls (FWs) are primarily designed for continuously connected environments and are therefore insufficient, by themselves, to address threats in physically isolated networks [44, 45]. Building on the comparative trait analysis in Section IV-C, air-gap defense should be grounded in recurring operational traits rather than in malware family names alone.

The recurring traits identified in Section IV-C provide the behavioral basis for the ACKC-aligned defenses in Table IV.

Unlike networked supply-chain security, which can rely on continuous vendor-portal monitoring and remote attestation, Closed Network Reconnaissance defenses must operate without external telemetry: supplier verification is instantiated as point-in-time physical audits during scheduled maintenance windows, and the listed honeypots are deployed as decoy engineering workstations or decoy PLC endpoints physically

TABLE VI
SCENARIO-BASED VALIDATION PROTOCOL FOR ACKC-ALIGNED DEFENSES

Validation Scenario	ACKC Phase	Observable Data Source	Defense to Evaluate	Primary Metrics
Controlled removable-media transfer	Malware Injection	USB serial number, mount event, file-copy log, scan result	USB allowlisting, media scanning, quarantine workflow	Block rate; scan latency; false-positive rate
Scripted file enumeration and staging	Stealth & Data Collection	File-access burst, archive creation, hidden directory, credential-use log	File-access anomaly detection, least privilege, decoy project files	Detection rate; MTTD; alert precision
Simulated ICS command anomaly	Attacks: Neutralization	PLC/SCADA command log, actuator state change, engineering workstation event	Logic restriction, command validation, execution isolation	Command-block rate; response time; recovery time
Controlled signal-emission experiment	Attacks: Data Exfiltration	RF/acoustic/optical emission pattern, spectrum-monitoring output	Emission shielding, spectrum monitoring, decoy emission, controlled jamming	Signal-detection rate; leakage-bandwidth reduction; false-alarm rate
Post-incident restoration drill	Attacks: post-incident restoration	Firmware integrity check, image restoration log, asset baseline comparison	System re-imaging, firmware restoration, clean-state verification	MTTR; restoration success rate; residual-risk indicator

present on the isolated network segment, whose access logs are reviewed locally rather than forwarded to an external security information and event management (SIEM) system.

For Malware Creation, defenders cannot ordinarily observe the external development process. A corresponding defensive opportunity arises later at intake: an isolated analysis station can check artifact integrity, inspect embedded trigger conditions, and emulate execution under controlled inputs. This is retrospective analysis of preparation artifacts, not prevention or observation of malware development itself. Physical or offline malware injection motivates media control, supplier verification, and removable-media auditing in the **Malware Injection** phase. Concretely, this is instantiated as a dedicated, disconnected media-scanning kiosk through which all removable media must pass before entering the protected segment, with scan results logged locally and reconciled against the work-order log described in Section IV-A.

Autonomous operation and stealthy data collection motivate execution restriction, file-access anomaly detection, least-privilege enforcement, deception, and persistence cleanup in the **Stealth and Data Collection** phase. Concretely, decoy project files are placed on engineering workstations with filenames matching real project naming conventions, and access to them is monitored via local file-system auditing rather than network-based data-loss-prevention tooling, which is unavailable in this environment.

Offline or covert exfiltration motivates emission monitoring, shielding, decoy emissions, and controlled jamming in the data-exfiltration branch of the **Attacks** phase. Concretely, this is instantiated as a fixed-position spectrum-monitoring sensor local to the protected segment (Section VI-D) rather than a cloud-connected SIEM feed, since no external network path is available to carry telemetry off-site.

Abnormal OT or hardware interaction motivates command validation, interface control, firmware-integrity checking, state restoration, and system re-imaging in the neutralization branch of the **Attacks** phase. Concretely, command validation is instantiated as a local PLC-logic allowlist enforced at the engineering workstation, since remote attestation services common in networked ICS deployments are not reachable from within the air gap.

Trigger or condition-aware logic is treated as a cross-phase mechanism rather than a standalone behavioral trait.




Fig. 11. Air-gap attack classification

It may be implemented during **Malware Creation**, remain dormant during **Stealth and Data Collection**, and manifest operationally during the **Attacks** phase when predefined time-, event-, or system-state conditions are satisfied. This distinction separates where the logic is designed, where it remains latent, and where its operational effect becomes observable.

These correlations characterize the operational logic of representative air-gap-capable malware, including Stuxnet, Havex, BlackEnergy3, and Industroyer2, and explain how the proposed defense matrix is derived. By treating recurring behavioral traits as ACKC phase indicators, defenders can implement phase-specific countermeasures, such as media isolation for injection, file and process monitoring for stealth, and emission shielding or restoration procedures for attack mitigation. Table VII summarizes representative malware characteristics, their usage examples, and their relevance to ACKC-based defense design.

D. Analysis of Air-Gap Attack Vectors and Defense Methods

As illustrated in Fig. 9, the **Attacks** phase of the ACKC includes covert data exfiltration through physical channels. Figure 11 organizes these channels into three operational groups. The electrical and electromagnetic group includes RF emissions, optical signaling, and power-line conduction. The mechanical group includes airborne acoustic signals and structure-borne vibration. The thermal group includes temperature-based signaling through heat transfer.

TABLE VII
AIR-GAP MALWARE CHARACTERISTICS AND ACKC-ALIGNED DEFENSE RELEVANCE

Feature	Description	Representative Usage Examples
Non-network path penetration	Malicious code infiltrates isolated environments via physical media or compromised supply chains.	Stuxnet and Havex used infected USB drives or malware-injected ICS software distributed through legitimate vendor channels [6, 46].
Propagation and lateral movement	Self-replication and credential-mediated lateral movement are distinct mechanisms; the latter does not by itself establish autonomous malware propagation.	Stuxnet used removable-media and local-network propagation mechanisms. The BlackEnergy3-associated campaign involved credential-mediated lateral movement; this is not evidence of autorun-based USB self-replication. [6, 37]
Stealth techniques	Use of common protocols and filenames to evade detection, combined with self-deletion and code obfuscation.	Havex disguised C&C communications using HTTP(S), BlackEnergy3 mimicked legitimate filenames, and Industroyer2 employed self-deletion [35, 38, 47].
Abnormal hardware communication	Exfiltration or disruption via manipulated hardware-level signals, abnormal OT commands, or process-control behavior.	Stuxnet altered PLC-controlled centrifuge speeds, while BlackEnergy3 used the KillDisk module to disrupt power grid operations [6, 37].

The following subsections discuss RF emissions and mechanical channels as representative signal-analysis examples, followed by thermal sensing and transfer constraints. Signal-analysis and mitigation methods require channel-specific adaptation and should not be assumed to apply unchanged across these groups.

1) *Electrical and Electromagnetic Channels*: RF electromagnetic channels exploit unintended or deliberately modulated radio-frequency emissions from electronic systems to establish covert communication with external receivers. As a specialized domain within Electromagnetic Compatibility (EMC), the TEMPEST discipline addresses both the analysis of unintentional emanations and the implementation of defensive countermeasures, such as zoning, shielding, and emission control.

This risk extends to industrial settings: Sachintha et al. demonstrated that electromagnetic (EM) radiation from wired Ethernet connections in ICS can be captured with a portable software-defined radio and demodulated to exfiltrate data, without requiring any direct physical connection to the target device [48].

In the context of air-gapped environments, adversaries can repurpose controllable electromagnetic emissions—originating from buses, cables, or peripheral interfaces—into data exfiltration pathways. Figure 12 illustrates a conceptual acquisition and analysis pipeline for RF electromagnetic emissions: received analog signals are digitized using an Analog-to-Digital Converter (ADC) and analyzed using a Fast Fourier Transform (FFT) to identify spectral components that may indicate covert modulation. This description is provided to clarify the defense rationale at the physical layer, rather than to prescribe site-specific implementation details. Once such patterns are detected, active suppression can be performed through destructive interference—by transmitting an inverse-phase waveform. For instance, under idealized single-frequency assumptions, if the emission is represented as $A \sin(\omega t)$, generating $A \sin(\omega t + \pi) = -A \sin(\omega t)$ illustrates the principle of phase cancellation. In practical deployments, however, the effectiveness of such active suppression depends on synchronization accuracy, propagation paths, background




Fig. 12. Conceptual FFT-based spectrum analysis of RF emissions.

noise, and receiver placement.

2) *Mechanical Channels*: Mechanical channels include airborne acoustic signals, such as speaker emissions or fan noise, and structure-borne vibration. The distinction concerns the propagation path and receiving sensor, rather than the emitting component alone.

Figure 13 illustrates conventional acoustic-frequency ranges: the infrasonic range below approximately 20 Hz, the audible range from approximately 20 Hz to 20 kHz, and the ultrasonic range above approximately 20 kHz. These approximate human-audibility boundaries do not define the usable bandwidth of a particular covert channel. The usable frequency range depends on the source, propagation path, and receiver. No universal upper frequency limit for the ultrasonic range is specified here, and the figure does not assign an operating frequency range to structure-borne vibration.

Doi and Sugawara demonstrate a covert channel in which CPU-load modulation induces leakage from motherboard ceramic capacitors, and a laser Doppler vibrometer is used instead of a microphone to recover the signal. This example highlights the need to consider the receiving measurement and optical access conditions when assessing acoustic or vibration-channel defenses [49].

Passive controls may include restricting unnecessary emission-capable components, acoustic insulation, and vibration damping, subject to equipment and cooling requirements.




Fig. 13. Conventional acoustic-frequency regions.

Active masking requires evaluation of the source, propagation path, receiver, and legitimate equipment operation.

SoniControl provides an implemented mobile-device example of ultrasonic-activity detection and user-controlled blocking. It uses a background spectral model for detection and supports acoustic jamming. Reported experiments also identify a limitation when the sender overwhelms the jamming signal. This provides implementation evidence for an acoustic defense, but does not establish detection or suppression performance in an isolated ICS deployment [50].

3) *Thermal Sensing and Transfer Constraints*: BitWhisper provides a temperature-based communication example using computer heat production and thermal sensing. Its evidence supports thermal-channel feasibility, not a general defense-detection benchmark [4].

A corresponding proposed defense would compare local temperature measurements with workload, fan state, ambient temperature, and maintenance activity. Sensor position, resolution, sampling interval, and thermal response time must be reported. Separation and airflow management can modify thermal coupling, but must preserve equipment cooling. RF spectrum monitoring or audio cancellation is not a substitute for this temperature-based evaluation.

VII. CONCLUSION

Despite extensive studies on individual air-gap cyber attacks, prior research has largely lacked a unified and systematized perspective that connects disparate attack types across their full operational life cycle. Defensive methodologies have often remained narrowly focused on specific covert channels, leaving critical phases such as malware infiltration, persistence, trigger-driven execution, and recovery comparatively underexplored.

To bridge this gap, this SoK systematizes existing knowledge through the ACKC, an air-gap-oriented adaptation of the CKC abstraction [9, 22]. The ACKC delineates sequential stages of air-gap attacks and provides a consistent analytical basis for reasoning about attacker progression and defender intervention points. By explicitly modeling physical or offline malware injection, constrained C&C, trigger-based execution, and signal-mediated operations, the ACKC captures operational constraints unique to air-gapped environments that

are not adequately represented in network-centric kill chain models [24, 26, 27].

Through comparative analysis of representative ICS malware families—including Stuxnet, Havex, BlackEnergy3, and Industroyer2 [6, 7, 34, 37]—we identified recurring behavioral traits relevant to closed-network operations and derived phase-specific defense indicators. We further examined recent air-gap-targeting campaigns, including Ramsay, USBCulprit, and PlugX/SOGU USB propagation variants [17, 18, 39, 40], to show that removable-media-mediated injection, autonomous operation, stealthy data collection, and delayed exfiltration remain relevant in modern air-gap scenarios. These findings informed the *Air-Gap Defense Matrix*, which aligns ACKC stages with observable indicators, defensive actions, and defense objectives to support situational awareness, early detection, and resilience against multi-stage air-gap threats.

Furthermore, physical covert channels were organized into three operational groups: electrical and electromagnetic, mechanical, and thermal. RF emissions and acoustic signals were examined as representative examples for detailed signal-based defense analysis [1, 3, 51–53]. The associated defenses were organized into passive sensing/monitoring-based detection and active countermeasure approaches. This multidimensional taxonomy highlights the need for integrated, layered defenses that combine physical safeguards, such as shielding, isolation, and noise injection, with logical controls, such as execution restriction, integrity verification, media control, and controlled maintenance workflows.

The channel examples reviewed in this SoK are not exhaustive. Hydraulic-pressure signaling through water-supply infrastructure illustrates a deployment-specific transfer path beyond the airborne acoustic and structure-borne vibration examples emphasized here. Such channels require explicit characterization of the pressure variable, infrastructure coupling, and receiving sensor when applying the operational taxonomy [54].

This work provides knowledge consolidation and an analytical framework, not a new empirical defense benchmark. The proposed reference configuration provides a path for future controlled evaluation, rather than evidence of operational effectiveness. Actual implementation, independently collected traces, and comparative performance measurements remain future work. The campaign corpus supports illustrative behavioral mapping and does not validate all attack vectors. Accordingly, the contribution is a metric-informed evaluation framework that may inform future standardization, not an established evaluation standard.

Future research should build on this validation protocol to empirically evaluate the ACKC and the Air-Gap Defense Matrix across diverse air-gap scenarios, including removable-media injection, stealthy data staging, simulated ICS command anomalies, controlled physical-channel emissions, and post-incident restoration drills. Such validation can enable quantitative assessment of detection coverage, false-alarm trade-offs, response effectiveness, operational cost, and residual risk. In parallel, integrating artificial intelligence (AI)-driven anomaly detection, cross-signal correlation, and automated response mechanisms may further enhance both passive and active

protection in real-world industrial and defense-grade deployments. Ultimately, this study aims to bridge structured threat modeling and practical defense design for physically isolated infrastructures operating in increasingly complex cyber-physical environments.

ACKNOWLEDGMENT

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (RS-2026-25480486) and the Institute of Information & Communications Technology Planning & Evaluation (IITP) under the virtual convergence support program to nurture the best talents (IITP-2026-RS-2023-00254529) grant funded by the Korea government (MSIT). The work of T. Taleb was supported in part by the Federal Ministry of Research, Technology, and Space (BMFTR), Germany, through the Project 6GEM+ under Grant 16KIS2411, and in part by the European Union's Horizon Europe research and innovation programme under the 6G-Path project (Grant No. 101139172).

REFERENCES

- [1] M. Guri, G. Kedma, A. Kachlon, and Y. Elovici, "Airhopper: Bridging the air-gap between isolated networks and mobile phones using radio frequencies," in *2014 9th International Conference on Malicious and Unwanted Software: The Americas (MALWARE)*. IEEE, 2014, pp. 58–67.
- [2] M. Guri, B. Zadov, D. Bykhovsky, and Y. Elovici, "Ctrl-alt-led: Leaking data from air-gapped computers via keyboard leds," in *2019 IEEE 43rd Annual Computer Software and Applications Conference (COMPSAC)*, vol. 1. IEEE, 2019, pp. 801–810.
- [3] M. Guri, Y. Solewicz, and Y. Elovici, "Fansmitter: Acoustic data exfiltration from air-gapped computers via fans noise," *Computers & Security*, vol. 91, p. 101721, 2020.
- [4] M. Guri, M. Monitz, Y. Mirski, and Y. Elovici, "Bitwhisper: Covert signaling channel between air-gapped computers using thermal manipulations," in *2015 IEEE 28th Computer Security Foundations Symposium*. IEEE, 2015, pp. 276–289.
- [5] M. Guri, B. Zadov, D. Bykhovsky, and Y. Elovici, "Powerhammer: Exfiltrating data from air-gapped computers through power lines," *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 1879–1890, 2019.
- [6] R. Team, "Cssh cyber defense project hotspot analysis: Stuxnet. no," 2017.
- [7] Forescout Research – Vedere Labs, "Industroyer2 and incontroller: In-depth technical analysis of the most recent ics-specific malware," Forescout Technologies, Tech. Rep., July 14 2022, published online – threat briefing. [Online]. Available: <https://www.forescout.com/resources/industroyer2-and-incontroller-report/>
- [8] Y. Ahmed, A. T. Asyhari, and M. A. Rahman, "A cyber kill chain approach for detecting advanced persistent threats," *Computers, Materials and Continua*, vol. 67, no. 2, pp. 2497–2513, 2021.
- [9] M. R. Na and K. Sundharakumar, "A study on air-gap networks," in *2024 5th International Conference on Innovative Trends in Information Technology (ICITIT)*. IEEE, 2024, pp. 1–6.
- [10] National Security Agency, "TEMPEST: A Signal Problem," <https://www.nsa.gov/portals/75/documents/news-features/declassified-documents/cryptologic-spectrum/tempest.pdf>, National Security Agency, Tech. Rep., 1972.
- [11] A. F. Gentile, P. Fazio, and G. Miceli, "A survey on the implementation and management of secure virtual private networks (vpns) and virtual lans (vlans) in static and mobile scenarios," in *Telecom*, vol. 2, no. 4. MDPI, 2021, pp. 430–445.
- [12] P. Chrobak, "Implementation of virtual desktop infrastructure in academic laboratories," in *2014 Federated Conference on Computer Science and Information Systems*. IEEE, 2014, pp. 1139–1146.
- [13] M. Guri, "Air-viber: Exfiltrating data from air-gapped computers via covert surface vibrations," *arXiv preprint arXiv:2004.06195*, 2020.
- [14] —, "Air-fi: Leaking data from air-gapped computers using wi-fi frequencies," *IEEE Transactions on Dependable and Secure Computing*, vol. 20, no. 3, pp. 2547–2564, 2022.
- [15] —, "Air-gap electromagnetic covert channel," *IEEE Transactions on Dependable and Secure Computing*, vol. 21, no. 4, pp. 2127–2144, 2023.
- [16] M. F. Bari and S. Sen, "Noisehopper: Emission hopping air-gap covert side channel with lower probability of detection," in *2024 IEEE International Symposium on Hardware Oriented Security and Trust (HOST)*. IEEE, 2024, pp. 21–32.
- [17] I. Sanmillan, "Ramsay: A cyber-espionage toolkit tailored for air-gapped networks," *Retrieved October*, vol. 12, p. 2020, 2020.
- [18] L. M. GReAT and G. Dedola, "Cycldek: Bridging the (air) gap," 2020.
- [19] Mandiant, "The spies who loved you: Infected usb drives to steal secrets," <https://cloud.google.com/blog/topics/threat-intelligence/infected-usb-steal-secrets/>, jul 2023, accessed: 2026-06-23.
- [20] E. H. Lo and T. A. Au, "Improving the kill chain for prosecution of time sensitive targets," *Dynamic Modelling*, vol. 290, 2010.
- [21] J. A. Tirpak, "Find, fix, track, target, engage, assess," *Air Force Magazine*, vol. 83, no. 7, pp. 24–29, 2000.
- [22] E. M. Hutchins, M. J. Cloppert, R. M. Amin *et al.*, "Intelligence-driven computer network defense informed by analysis of adversary campaigns and intrusion kill chains," *Leading Issues in Information Warfare & Security Research*, vol. 1, no. 1, p. 80, 2011.
- [23] N. Kollasch, "Data loss prevention and a point of sales breach," SANS Institute, Global Information Assurance Certification (GIAC) Gold, USA, GIAC (GSEC) Gold Certification Paper, 2014, accepted: July 12, 2014.

- [Online]. Available: <http://www.giac.org>
- [24] H. P. Enterprise, “Hpe attack life cycle use case methodology,” *Technical White Paper*, 2016.
- [25] T. Grant, “Detect, deny, degrade, disrupt, destroy, deceive: Which is the greatest in oco?” in *ECCWS 2023 22nd European Conference on Cyber Warfare and Security*, no. 1. Academic Conferences and publishing limited, 2023.
- [26] O. Alexander, M. Belisle, and J. Steele, “Mitre att&ck for industrial control systems: Design and philosophy,” *The MITRE Corporation: Bedford, MA, USA*, vol. 29, 2020.
- [27] B. I. Messaoud, K. Guennoun, M. Wahbi, and M. Sadik, “Advanced persistent threat: New analysis driven by life cycle phases and their challenges,” in *2016 International conference on advanced communication systems and information security (ACOSIS)*. IEEE, 2016, pp. 1–6.
- [28] R. Kour, R. Karim, and P. Dersin, “Modelling cybersecurity strategies with game theory and cyber kill chain,” *International Journal of System Assurance Engineering and Management*, pp. 1–12, 2025.
- [29] M. Kim, J.-H. Cho, H. Lim, T. Moghaddam, T. J. Moore, F. F. Nelson, and D. D. Kim, “Mtd in depth: Multi-phased moving target defense techniques against cyber-attacks based on cyber kill chain,” *Future Generation Computer Systems*, vol. 181, p. 108419, 2026.
- [30] A. Presekal, A. Ștefanov, V. S. Rajkumar, I. Semertzis, and P. Palensky, “Advanced persistent threat kill chain for cyber-physical power systems,” *IEEE Access*, vol. 12, pp. 177 746–177 771, 2024.
- [31] J. Park, J. Yoo, J. Yu, J. Lee, and J. Song, “A survey on air-gap attacks: Fundamentals, transport means, attack scenarios and challenges,” *Sensors*, vol. 23, no. 6, p. 3215, 2023.
- [32] Y. Mekdad, G. Bernieri, M. Conti, and A. El Fergougui, “The rise of ics malware: A comparative analysis,” in *European Symposium on Research in Computer Security*. Springer, 2021, pp. 496–511.
- [33] MITRE Corporation, “Common Vulnerabilities and Exposures (CVE),” 2025. [Online]. Available: <https://cve.mitre.org/>
- [34] M. Hemsley and E. Fisher, “History of industrial control system cyber incidents,” Idaho National Laboratory, Tech. Rep. INL/EXT-18-44708-Revision-0, 2018, accessed: 2025-04-27. [Online]. Available: <https://www.osti.gov/servlets/purl/1505628>
- [35] E. Hjeltnvik, “Observing the havex rat,” <https://www.netresec.com/?month=2014-11&page=Blog&post=Observing-the-Havex-RAT>, November 2014, accessed: 2025-05-26.
- [36] J. Krushi, H. Farhangi, C. Howey, K. Carmichael, and J. Dabell, “A quantitative evaluation of the target selection of havex ics malware plugin,” in *Industrial control system security (ICSS) workshop*, 2015, pp. 1–5.
- [37] D. U. Case, “Analysis of the cyber attack on the ukrainian power grid,” *Electricity information sharing and analysis center (E-ISAC)*, vol. 388, no. 1-29, p. 3, 2016.
- [38] A. Cherepanov and R. Lipovsky, “Industroyer2: Industroyer reloaded,” <https://www.welivesecurity.com/2022/04/12/industroyer2-industroyer-reloaded/>, 2022, accessed: 2025-05-12.
- [39] M. Guri, “Usbculprit: Usb-borne air-gap malware,” in *Proceedings of the 2021 European Interdisciplinary Cybersecurity Conference*, 2021, pp. 7–13.
- [40] MITRE ATT&CK, “Plugx, software s0013,” <https://attack.mitre.org/software/S0013/>, 2026, accessed: 2026-06-23.
- [41] M. Guri, “Mind the gap: Can air-gaps keep your private data secure?” *arXiv preprint arXiv:2409.04190*, 2024.
- [42] M. Guri, A. Kachlon, O. Hasson, G. Kedma, Y. Mirsky, and Y. Elovici, “{GSMem}: Data exfiltration from {Air-Gapped} computers over {GSM} frequencies,” in *24th USENIX Security Symposium (USENIX Security 15)*, 2015, pp. 849–864.
- [43] M. Guri, Y. Solewicz, and Y. Elovici, “Mosquito: Covert ultrasonic transmissions between two air-gapped computers using speaker-to-speaker communication,” in *2018 IEEE Conference on Dependable and Secure Computing (DSC)*. IEEE, 2018, pp. 1–8.
- [44] A. Trisolino, “Analysis of security configuration for ids/ips,” Ph.D. dissertation, Politecnico di Torino, 2023.
- [45] K. Neupane, R. Haddad, and L. Chen, “Next generation firewall for network security: a survey,” in *SoutheastCon 2018*. IEEE, 2018, pp. 1–6.
- [46] C. Clancy, J. Ferraro, R. Martin, A. Pennington, C. Sledjeski, and C. Wiener, “Deliver uncompromised: Securing critical software supply chains,” *MITRE Technical Papers*, vol. 24, no. 01, 2021.
- [47] GREAT (Kaspersky Lab), “Blackenergy apt attacks in ukraine employ spear-phishing with word documents,” <https://securelist.com/blackenergy-apt-attacks-in-ukraine-employ-spearphishing-with-word-documents/73440/>, jan 2016, accessed: 2025-08-10.
- [48] S. Sachintha, N.-A. Le-Khac, M. Scanlon, and A. P. Sayakkara, “Data exfiltration through electromagnetic covert channel of wired industrial control systems,” *Applied Sciences*, vol. 13, no. 5, p. 2928, 2023.
- [49] K. Doi and T. Sugawara, “Laser-based covert channel attack using inaudible acoustic leakage from multilayer ceramic capacitors,” *IEICE TRANSACTIONS on Fundamentals of Electronics, Communications and Computer Sciences*, vol. 108, no. 3, pp. 207–214, 2025.
- [50] M. Zeppelzauer, A. Ringot, and F. Taurer, “Sonicontrol-a mobile ultrasonic firewall,” in *Proceedings of the 26th ACM international conference on Multimedia*, 2018, pp. 1250–1252.
- [51] A.-M. Vizitiu, L. Dobrescu, C. Molder, B. Sebacher, B. C. Trip, and V. F. Butnariu, “Detection of tempest audio compromising signal using artificial intelligence,” in *2024 15th International Conference on Communications (COMM)*. IEEE, 2024, pp. 1–6.
- [52] A. Auddy and S. Sahu, “Tempest: Magnitude of threat and mitigation techniques,” in *2008 10th International Conference on Electromagnetic Interference & Compatibility*. IEEE, 2008, pp. 603–611.
- [53] Interelectronix, “Us and nato tempest standards,” n.d.

[Online]. Available: <https://www.interelectronix.com/tempest.html>

- [54] Y. Liang, H. Shan, J. Zhao, C. Lu, and G. Yang, "Hydraulicbridge: Covert signaling channel between air-gapped systems using hydraulic-pressure fluctuations," *Electronics*, vol. 13, no. 15, p. 3010, 2024.



Bonn Park (Member, IEEE) received the B.S. degrees in defense systems engineering and electronics, information and communications engineering (double major) from Sejong University, Seoul, South Korea, in 2019, and the M.S. degree in information security from the Graduate School of Information Security, Sejong Cyber University, Seoul, South Korea, in 2022. He is currently pursuing the Ph.D. degree in information security with the Department of Artificial Intelligence Convergence, Sejong University, having completed the coursework. He served

for seven years as an information and communications officer in the Republic of Korea Navy, where he was responsible for information security operations, and was discharged with the rank of Lieutenant. His research interests include closed-network security, malware analysis, and security planning.



JiEun Lee (Member, IEEE) received the Ph.D. degree from Sejong University, Seoul, South Korea. During her Ph.D. studies, she spent one year as an invited researcher with Easy Global Market (EGM), France, where she participated in ETSI MEC, PDL, and CIM activities. She was previously with the Korea Electronics Technology Institute (KETI). She is currently a Research Professor with the Department of Intelligence and Information Convergence, Sejong University, Seoul, South Korea. She has contributed to oneM2M and ETSI MEC standardization. Her

research interests include standards-based IoT platforms, IoT and edge computing interworking, and semantic interoperability.



JiHo Lee (Member, IEEE) received the B.S. and M.S. degrees in Computer and Information Security from Sejong University, Seoul, South Korea, in 2023 and 2025, respectively, and is currently pursuing the Ph.D. degree. From 2023 to 2025, he was a blockchain developer with Guarantable. From 2023 to 2025, he served as Rapporteur for the ETSI ISG Permissioned Distributed Ledger specification PDL-028. His research interests include permissioned blockchains, standards, MCP, and physical AI.



Tarik Taleb (Senior Member, IEEE) received the B.E. degree (with distinction) in information engineering and the M.Sc. and Ph.D. degrees in information sciences from Tohoku University, Sendai, Japan, in 2001, 2003, and 2005, respectively. He is currently a Full Professor of Networked Energy Efficient Systems with the Faculty of Electrical Engineering and Information Technology, Ruhr University Bochum, Bochum, Germany. He was a Professor with the Center for Wireless Communications, University of Oulu, Oulu, Finland. From October

2014 to December 2021, he was an Associate Professor with the School of Electrical Engineering, Aalto University, Espoo, Finland. Prior to that, he was a Senior Researcher and a 3GPP Standards Expert with NEC Europe Ltd., Heidelberg, Germany. Until March 2009, he was an Assistant Professor with the Graduate School of Information Sciences, Tohoku University. He is the founder of ICTFICIAL Oy and the founder and Director of the MOSA!C Lab, Espoo, Finland.



JaeSeung Song (Senior Member, IEEE) received the B.S. and M.S. degrees in computer science from Sogang University, Seoul, South Korea, and the Ph.D. degree from Imperial College London, London, U.K. Prior to joining Sejong University, he held positions at NEC Europe Ltd. and LG Electronics. He is currently a Distinguished Professor with the Department of Computer and Information Security, Sejong University, Seoul, South Korea. He serves as the Technical Plenary Vice Chair of the oneM2M global IoT standards initiative. His

research interests include beyond 5G/6G, AI/ML-enabled network systems, software engineering, and security for physical AI systems.