CONTRACT-CONSTRAINED CROSS-LAYER COORDINATION FOR CYBER RESILIENCE

Oleksandr Korchenko (1), Dmytro Prokopovych-Tkachenko (2), Svitlana Kazmirchuk (1), Mukafat Askerov (1), Iurii Zhurov (3)*

(1) State University of Information and Communication Technologies (SUICT), Kyiv; (2) University of Customs and Finance, Dnipro; (3) SMB Cybersecurity Advisors LLC, Palm Harbor, Florida
(1,2) Ukraine, (3) USA

* Corresponding Author, e-mail: s.kazmirchuk@duikt.edu.ua

Abstract: Locally correct cyber-protection responses can amplify disruption across communication layers. This study specifies a contract-constrained coordination framework that evaluates protection and recovery against mission quality. The method combines a seven-layer reference mapping, qualified evidence, causal dependencies, robust state estimation, normalized resilience measures, and safety and stability guards. A worked interference cascade demonstrates the decision sequence. Bounded explicit-state exploration of the contract lifecycle covers 1,348 reachable states and 2,320 transitions, with six properties satisfied under the full guard set. Removing duplicate detection, conflict precedence, or forced expiration produces counterexamples. The contribution is a traceable coordination specification with initial bounded formal evidence; empirical benefits remain untested and require simulation, hardware-in-the-loop evaluation, pilot deployment, and independent replication.

Key words: cyber resilience, cross-layer coordination, action contracts, trusted telemetry, formal verification.

1. INTRODUCTION

Modern mission-critical information and communication systems span physical, link, network, transport, session, presentation, and application functions that are engineered and defended largely in isolation. Protection and recovery mechanisms at each layer are usually correct with respect to their own objectives, such as restoring coverage or reachability, limiting stalled connections, or removing a suspicious identity. When such mechanisms act concurrently on the same disturbance, however, their combined effect on the mission is not guaranteed to be beneficial. A channel change can invalidate a route that is being repaired, a revocation can lock out a controller that is executing recovery, and simultaneous rate and route adaptations can oscillate. Cross-layer disruption of this kind can turn locally correct defences into mission loss.

Two developments make the problem urgent. First, programmable infrastructure (software-defined networking, programmable data planes, zero-trust access control, and network digital twins) has made automated cross-layer actuation technically feasible, so the cost of an uncoordinated automated response is now measured in seconds rather than in operator shifts. Second, adversaries increasingly target the observation channel itself: telemetry can be replayed, suppressed, poisoned, or rolled back, so a coordinator that acts on unqualified evidence can be steered into harmful actions by the very attacker it is meant to counter. Coordination therefore has to be conditioned on the quality of the evidence, on the causal dependencies between layers, on explicit authority, and on stability bounds.

The scientific problem is the lack of a formal, checkable rule for admitting and ordering protective actions that different layers issue against one disturbance. Resilience assessment measures loss but does not govern the response; cross-layer detection acts within one domain without bounding action lifetime, budget, or reversibility; and zero-trust authority decides who may act, not whether an action is premature or destabilizing.

The contribution is a traceable specification with initial bounded formal evidence, not a validated system. Theoretically, it addresses each of these gaps, which Section 2 traces in the individual works: mission quality (5) becomes the objective of action selection (8), and the capped resilience (7) becomes the criterion for comparing policies; every cross-layer action is issued as a contract (10) that bounds lifetime, budget, and reversibility and records provenance and authorization; and the admissibility set (9) adds safety and stability conditions to zero-trust authorization. In addition, unlike centralized optimization, which assumes trustworthy observations, evidence qualification (2)–(4) discounts replayed, stale, or inconsistent records. Bounded exploration with guard ablations (Table 7) supports the contract rules, and Table 5 compares the resulting scope with seven approach classes. Practically, it offers application scenarios only, namely a worked cascade and a validation ladder; no accuracy, latency, overhead, operational gain, or superiority is claimed.

2. RELATED WORK AND MOTIVATION

Cyber-resilience research has moved from qualitative principles toward assessment frameworks and measurable indicators. A systematic review of cyber-resilience assessment frameworks [1] finds that most frameworks describe capabilities and maturity levels but rarely connect them to time-resolved mission performance. The recent survey of cyber-resilience strategies [2] identifies continuity, recovery, and adaptation as recurring objectives and stresses the absence of a common quantification. Availability-based quantification with area-under-curve normalization [3] provides a time-integrated loss measure that the resilience measure (7) below follows in spirit while adding a preapproved reference trajectory and a loss cap. None of these works, however, addresses how protective actions at different layers should be sequenced so that the measured resilience is not degraded by the response itself.

Cross-layer design has a long record of warning about unstable coupling. The parametric survey of cross-layer designs for wireless sensor networks [4] catalogues interactions between physical-, link-, and network-layer adaptations and shows that

control loops formed across layers must be analysed for stability. Explainable cross-layer machine learning for enterprise WLAN auditing [5] demonstrates that combining physical- and link-layer indicators with interpretable models improves the traceability of security findings, the kind of qualified evidence a coordinator needs. These studies motivate the explicit dependency structure and the dwell and hysteresis constraints adopted below.

Programmable infrastructure provides the actuation substrate. The survey of security applications of P4 programmable switches [6] and the survey of data-plane programming with P4 [7] document fine-grained, line-rate enforcement that can implement admission, filtering, rerouting, and isolation within microseconds; the same capability makes an erroneous or replayed control action equally fast. Zero-trust architecture [8] constrains who may act and under which continuously re-evaluated conditions, and it is adopted here as the authority model behind the Auth predicate. A survey of network digital twin architectures for the edge–cloud continuum [9] and a systematic review of network digital twins [10] identify traceable, synchronized, and versioned models as prerequisites; the versioned knowledge base $\mathcal{K}$ and the provenance field of the contracts serve this requirement.

Existing cross-layer defences concentrate on detection and on optimal response within a single domain. Data-driven cross-layer real-time attack detection in 5G Open RAN [11] and cross-layer detection and defence against DDoS and DRDoS attacks with P4 switches [12] show that fusing indicators from several layers improves detection, and automated cross-layer security for zero-touch 6G networks [13] extends this to closed-loop response. Spatial localization of Wi-Fi threat sources through correlation of radio events with security-operations events [14] illustrates how radio-layer evidence enters an operations workflow and why such evidence must carry uncertainty and provenance. Declarative secure placement of orchestrations in the cloud–edge continuum [15] treats security requirements as explicit constraints on placement, which the admissibility set (9) generalizes to arbitrary actions. Taken together, these approaches address complementary but not composable objectives: none of them specifies how detection outputs, authority, budgets, and rollback are combined into a mission-level decision across layers. Methodologically, layer-isolated control cannot prevent cross-layer interference, uncoordinated concurrent response issues actions without shared precedence, budgets, or lifetimes, and centralized optimization assumes trustworthy observations and leaves reversal unbounded; the proposed method keeps the global objective but adds qualified evidence (3)–(4) and contracts (10), and these three families serve as baselines in Section 4.4.

We use the Open Systems Interconnection (OSI) reference model (ISO/IEC 7498-1) only as a coordinate system for evidence and actions, and we take cyber-resilient systems engineering [16], zero-trust architecture [17], and the Cybersecurity Framework 2.0 [18] as implementation constraints rather than as methods. The reviewed corpus is purposive: it comprises fifteen research publications, thirteen of which date from 2022–2026, and three guidance documents. This requirement synthesis leads to the hypothesis examined in the remainder of the paper: coordination benefits are conditioned on qualified evidence, causal ordering, explicit authority, and stability bounds, and they must be demonstrated against declared baselines rather than assumed.

3. PROBLEM AND PROPOSED COORDINATION METHOD

3.1. System model and assumptions

Mission cyber resilience preserves, restores, or safely degrades prioritized functions under faults, attacks, or deceptive observations. Adaptive coordination governs authorized actions across at least two layers. Signed telemetry may be false; partitions require safe local policies; effects are delayed; adversaries can replay, suppress, poison, or roll back messages. The model includes layer-specific attacks and correlated faults, excluding physical destruction, undisclosed cryptographic breaks, and compromise of every independent source.

The domain is represented by the ten-component object

$$\mathcal{M}_{\mathrm{OSI}}(t) = \langle \mathcal{L}, \mathcal{V}, \mathcal{E}_x, s(t), y(t), \mathcal{A}, \mathcal{P}, \Theta, \mathcal{R}, \mathcal{K} \rangle, \quad (1)$$

Here $\boldsymbol{\mathcal{L}}$ = {L1, …, L7} denotes reference layers; $\boldsymbol{\mathcal{V}}$ entities; $\boldsymbol{\mathcal{E}}_x$ temporal dependencies; s(t) latent state; y(t) observations; $\boldsymbol{\mathcal{A}}$ actions; $\boldsymbol{\mathcal{P}}$ policy and authority; Θ thresholds and hysteresis; $\boldsymbol{\mathcal{R}}$ resource budgets; and $\boldsymbol{\mathcal{K}}$ versioned causal, topology, protocol, and mission knowledge. Each dependency e = (u, v, ℓ, δ, σ) specifies source, destination, type, delay, and influence, with provenance, applicability, and uncertainty. Correlation alone does not establish propagation. Table 1 maps these roles to concrete mechanisms.

Table 1. Reference-layer evidence, disturbances and bounded responses.

Layer	***Evidence/threat → bounded action; mission effect***
L7 App.	*Deadlines/denial → admission, priorities, replicas; completion.*
L6 Pres.	*Crypto/key misuse → rekey, policy, quarantine; security–continuity trade-off.*
L5 Session	*Identity/replay → revoke, reauthenticate, fallback; reachability/access.*
L4 Trans.	*Loss/flood → rate, connection, multipath, timeout controls; timeliness.*
L3 Net.	*Paths/poisoning → reroute, isolate, segment, diversity; reachability/load.*
L2 Link	*Retries/spoofing → handover, queues, filters, access control; delay/coverage.*
L1 Phys.	*Signal/jamming → channel, power, coding, redundancy; coverage/energy.*

Layers are reference coordinates; mappings are many-to-many.

3.2. Evidence qualification and estimation

Records identify source, acquisition/receipt times, sequence/epoch, units, calibration, entities/layers, policy version, and uncertainty. Their quality vector is

$$\mathbf{z}_r(t) = [I_r(t), P_r(t), F_r(t), C_r(t), D_r(t), U_r(t)]^{\mathsf{T}}, \quad (2)$$

The components I, P, F, C, D, U lie in [0,1] and denote integrity, provenance, freshness, consistency, diversity, and uncertainty (Table 2). Larger U means weaker evidence. Let $d_r = [I_r, P_r, F_r, C_r, D_r, 1 - U_r]$. The conservative trust weight and state estimate are

$$\tau_r(t) = \prod_{k=1}^{6} d_{r,k}(t)^{\alpha_k}, \quad \alpha_k > 0, \quad \sum_{k=1}^{6} \alpha_k = 1, \quad (3)$$

$$\hat{s}(t) = \underset{s \in \mathcal{S}}{\operatorname{argmin}} \left[\sum_{r \in \mathcal{O}(t)} \tau_r(t) \rho_r\big(y_r(t) - h_r(s)\big) + \lambda \Omega(s) \right], \quad (4)$$

Positive α_k prevent compensation for zero components; integrity/freshness floors quarantine unsuitable evidence. In (4), $\boldsymbol{\mathcal{S}}$ is the state space, $\boldsymbol{\mathcal{O}}(t)$ qualified records, h_r the observation map, ρ_r a nonnegative robust loss, Ω a nonnegative consistency regularizer, and $\lambda \geq 0$ its weight. Ties use a declared deterministic rule; weights require calibration.

Table 2. Evidence-quality dimensions and failure responses.

Dimension	***Qualification → failure response***
I: Integrity	*Signature/message authentication/channel → reject, quarantine, reduce weight.*
P: Provenance	*Source/calibration/transform chain → flag incompleteness; block high-impact action.*
F: Freshness	*Times/sequence/clock/TTL → expire, refresh, degraded mode.*
C: Consistency	*Residual/topology/protocol/mission agreement → robust estimate/conflict workflow.*
D: Diversity	*Independent sensors/collectors/clocks/trust roots → discount correlated evidence.*
U: Uncertainty	*Calibration/timing/model bounds → widen guards or require approval.*

3.3. Mission quality and bounded resilience

For $J \geq 1$ approved attribute, $q_j(t)$ may represent timeliness, delivery, confidentiality, integrity, reachability, or safety. Mission quality, capped loss, and resilience over $T > 0$ are

$$Q_m(t) = \sum_{j=1}^{J} \omega_j \, q_j(t), \quad \omega_j \geq 0, \quad \sum_{j=1}^{J} \omega_j = 1, \quad 0 \leq q_j(t) \leq 1, \tag{5}$$

$$L_m(t) = \min\{L_{\max}, \max[0, Q_m^{\mathrm{ref}}(t) - Q_m(t)]\}, \tag{6}$$

$$R_m(T) = 1 - \frac{\int_0^T L_m(t)\,\mathrm{d}t}{T L_{\max}}. \tag{7}$$

Mission owners preapprove ω_j, reference $Q_m^{\mathrm{ref}}(t)$, and cap $0 < L_{\max} \leq 1$. Critical attributes impose hard constraints. Comparison requires compatible missions/horizons; recovery objectives, tolerated outages, and safe degradation remain explicit policies.

3.4. Action selection and contracts

Actions are atomic or ordered plans. Over $H > 0$, selection follows

$$a^*(t) = \underset{a \in \mathcal{A}_{\mathrm{adm}}(t)}{\operatorname{argmax}} \{\mathbb{E}[Q_m(t+H) \mid \hat{s}(t), a] - \lambda_c C(a) - \lambda_r \mathrm{Risk}(a) - \lambda_o \mathrm{Osc}(a)\} \tag{8}$$

The expectation uses a versioned model and declared disturbance distribution. C, Risk, and Osc denote nonnegative normalized cost, residual risk, and reversal propensity; λ coefficients are nonnegative mission-approved weights. Safety-first precedence, then cost, resolves ties. Admissibility requires

$$\mathcal{A}_{\mathrm{adm}}(t) = \{a \in \mathcal{A} : \mathbf{g}_{\mathrm{safe}}(a, \hat{s}) \preccurlyeq \mathbf{0} \wedge \mathrm{Auth}(a) = 1 \wedge C(a) \leq B(t) \wedge \mathbf{g}_{\mathrm{stab}}(a, \Theta) \preccurlyeq \mathbf{0}\} \tag{9}$$

Safety/stability inequalities are componentwise; Auth checks authority. B(t) derives from $\boldsymbol{\mathcal{R}}$ on the scale of C. Rate, dwell, hysteresis, causal-order, and hold-down constraints govern stability. No admissible action means degradation, containment, or escalation. Cross-domain contracts follow (10) and Table 3:

$$\mathcal{C}_{ij} = \langle \mathrm{pre}, \mathrm{act}, \mathrm{post}, B, \mathrm{TTL}, \mathrm{rb}, \mathrm{prov}, \mathrm{auth} \rangle, \tag{10}$$

Table 3. Action-contract fields, enforcement and controlled risks.

Field	***Content; enforcer → controlled risk***
pre	*State, evidence floors, causal prerequisites, mission protection; selector/adapter → deceptive/stale evidence.*
act	*Typed operation, target, bounds, idempotency; adapter → ambiguity/duplication.*
post	*Expected effect, window, success/failure predicates; monitor/estimator → premature success.*
B	*Bandwidth/compute/energy/risk/attention reservation; resource manager → defensive exhaustion.*
TTL	*Activation, lifetime, epoch, renewal; adapter/reconciler → stale controls/replay.*
rb	*Rollback state, trigger, fallback, dependency order; recovery orchestrator → unsafe reversion.*
prov	*Evidence/model/policy/code/configuration versions; audit ledger → irreproducibility.*
auth	*Approver, role, scope, separation of duty; authorization service → privilege/conflict risk.*

Issuing/target domains i and j enforce scope and lifetime independently. Actions are idempotent or duplicate-checked; conflicts follow safety, mission, authority, and causality. Traces preserve evidence/model versions, rejected actions, predicted/observed effects, and completion/rollback.

Coordination gain compares policies under identical disturbance ξ, initial state s_0, and cumulative resource envelope $\boldsymbol{\mathcal{B}}$:

$$\Delta_{\mathrm{em}}(m, T \mid \xi, s_0, \mathcal{B}) = \mathbb{E}[R_m^{\mathrm{coord}}(T) \mid \xi, s_0, \mathcal{B}] - \mathbb{E}[R_m^{\mathrm{local}}(T) \mid \xi, s_0, \mathcal{B}], \quad (11)$$

Both terms use (7), with expectations over registered observation noise, estimator uncertainty, and randomized policy. The shared $\boldsymbol{\mathcal{B}}$ derives from $\boldsymbol{\mathcal{R}}$ and bounds cumulative costs/budgets. Positive gain needs a declared baseline and uncertainty; causation requires controlled comparison. For N actions in T, unintended reversal is

$$O(T) = \frac{1}{N-1} \sum_{k=2}^{N} \mathbf{1}\,[a_k \sim \mathrm{inv}(a_{k-1})], \quad N \geq 2. \quad (12)$$

The action catalogue defines the inverse relation. Planned contract-compliant rollback is excluded, and O(T) = 0 for N < 2. Dwell, hysteresis, and settling-time parameters require calibration.

4. RESULTS AND DISCUSSION

4.1. Architecture and analytical properties

Figure 1 traces one pass of the coordination loop. The trusted evidence plane scores telemetry from layers L1–L7 by (2)–(3) and quarantines records below the integrity or freshness floors. Qualified evidence then feeds the robust state estimator (4), the causal model of inter-layer dependencies, and the mission-quality monitor (5)–(7), which also receives application outcomes over the dashed mission-feedback path. Constrained policy selection (8) ranks candidate actions by predicted effect; the chosen action must pass the safety, authorization, and stability gate (9) and is issued as a contract (10), whose bounded controls return to the layers and whose effects re-enter as telemetry. Degraded

autonomy restricts cached policies, budgets, and lifetimes; recovery reconciles epochs and invalidates stale contracts.

Proposition 1: integrating $0 \le L_m(t) \le L_{max}$ gives $0 \le \int_0^T L_m(t)dt \le TL_{max}$, hence $0 \le R_m(T) \le 1$. Proposition 2: positive Δ_{em} indicates expected baseline advantage, but unequal resources, model error, or bias can confound causation. Proposition 3: admissible selection plus independent scope/TTL enforcement preserves represented invariants only under a sound safety model; unknown hazards and target compromise remain excluded.

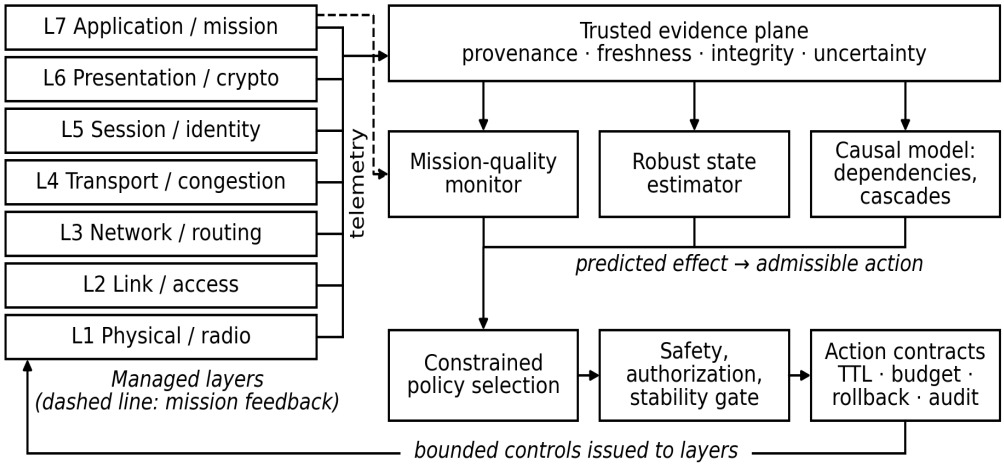


Figure 1. Coordination architecture: layer telemetry is qualified, estimated, and causally assessed before the safety, authorization, and stability gate issues bounded action contracts.

4.2. Worked interference cascade

Figure 2 and Table 4 trace interference to missed deadlines. Coordination orders reversible actions by qualified evidence and causal proximity; concurrent channel, route, timeout, and identity changes can amplify loss.

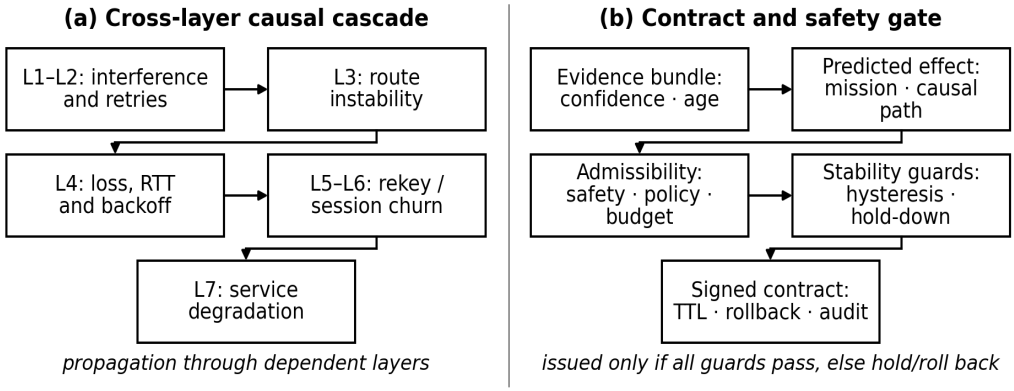


Figure 2. Worked cascade: (a) propagation from L1–L2 interference to L7 degradation; (b) guards that must all pass before a contract is issued.

Table 4. Interference-to-mission cascade: evidence, response and guards.

Stage	***Evidence → response; guard and mission effect***
Physical/link	*Independent signal/spectrum/error evidence → channel/power/coding/handover; freshness, protected-channel evidence, energy/coverage budgets, TTL; preserve coverage.*
Queue growth	*Queues/access/retries → priorities/shedding; mission class, loss bound, rollback after stabilization; protect priority traffic.*
Route reaction	*Epoch/path/topology → hold/reroute/multipath; hysteresis, hold-down, spare capacity, isolation authority; avoid flapping/overload.*
Transport	*RTT/loss/connections/deadline → rate/timeout/protection; minimum goodput, fairness, route dependency; preserve deadlines.*
Identity/session	*Token/key/anomaly → rekey/revoke/grace/local credential; zero trust, separation of duty, safety exception, short TTL; avoid lockout.*
Application	*Mission quality/dependencies → admission/priority/replica/degradation; safety floor, budget, escalation; restore or safely degrade mission.*

4.3. Positioning against existing approaches

Unlike the other tables, which present the authors' method, scenario, validation plan, and results, Table 5 refers to other works. It is not a review: these works are discussed in Section 2 and serve here only as coded representatives of approach classes, so that the contribution can be positioned against them.

Table 5. Positioning of the proposed framework against approach classes (0 = not central; 1 = partial/implicit; 2 = explicit; SDN = software-defined networking).

Approach class	***Coded works***	***(D1,D2,D3,D4,D5)***	***Σ /10***
CLS: Cross-layer surveys	*[4], [5]*	*(1,1,1,1,0)*	*4*
CRM: Resilience assessment/metrics	*[1], [2], [3]*	*(1,2,0,0,1)*	*4*
ZT: Zero trust (individually coded)	*[8]*	*(1,0,1,1,1)*	*4*
SDN: SDN/programmable data planes	*[6], [7]*	*(1,1,2,1,2)*	*7*
NDT: Network digital twins/security	*[9], [10]*	*(2,1,1,1,2)*	*7*
CLD: Cross-layer detection/defense	*[11], [12], [13], [14]*	*(2,1,2,1,2)*	*8*
CEI: Secure cloud–edge placement	*[15]*	*(2,1,2,2,1)*	*8*
Proposed framework	—	*(2,2,2,2,1)*	*9*

D1–D5 denote evidence/causality (2)–(4), mission trajectory (5)–(7), contracts (10), safety/stability (9), and deployment/validation. Only the proposed framework is explicit on D2–D4 together; its D5 stays partial until V1–V5 are completed. Scores describe scope, not superiority; independent recoding and inter-rater agreement remain pending.

4.4. Validation ladder and bounded lifecycle results

Table 6 separates bounded V0 from planned V1–V5. Compare layer-isolated, uncoordinated concurrent, and centralized uncontrolled baselines under matched seeds, topology, workload, budgets, and missions. Ablate evidence qualification, causal prediction, mission feedback, and stability guards separately.

Table 6. Validation levels, comparisons and minimum exit evidence.

V	***Question/method; comparison, measures and exit evidence***
V0	*Internal safety/authorization/liveness/stability: model checking with adversarial states; no unsafe reachability, deadlock, unauthorized transition or unbounded reversal within declared scope; assumptions/gaps disclosed.*
V1	*Mission benefit: attack/fault simulation against isolated/uncoordinated/ablated baselines; R_m, Δ_{em}, false actions, recovery, O(T); effects and uncertainty across seeds/loads/topologies/adversaries.*
V2	*Hardware/protocol fidelity: hardware-in-the-loop with V1 scenarios; latency, resources, clock error, rollback; bounded deadlines/budgets, no unmodeled unsafe actuation.*
V3	*Operational feasibility: approved pilot/red team; current-procedure/V2 baselines; service/operator/audit measures; mission/safety criteria met, incidents documented.*
V4	*Transferability: heterogeneous replication; local baselines/common protocol; disclose heterogeneity, recalibration cost and failure limits.*
V5	*Reproducibility: independent team, released artifacts, adverse cases; registered hypotheses/common schema; independent result, deviations, negative findings, maturity.*

Initial V0 uses exhaustive breadth-first exploration of a Python automaton: two conflicting contracts, six clock ticks, TTL two ticks, activation budget three, and dwell one tick. States are proposed, authorized, active, completed, expired, rolled back, and revoked. Adversaries replay consumed activations; guards enforce authorization, duplicate detection, conflict precedence, dwell, and expiration. The environment is a self-contained explicit-state explorer in Python whose state records each contract phase, the clock, the remaining budget, the dwell status, and consumed activations. The parameters are the smallest that expose each hazard: two contracts create a conflict, a two-tick TTL within six ticks admits activation, expiry, and post-expiry replay in one run, a budget of three leaves room for a replay, and a one-tick dwell makes reversal observable. Each ablation removes one guard against one threat of Section 3.1. The script is available from the corresponding author on request.

Properties are P1 prior authorization; P2 no second activation by replay; P3 no activity past TTL after clock advance; P4 no post-expiry reactivation and only declared terminal exits; P5 no conflicting active contracts; P6 dwell-bounded reversals. Table 7 reports full-guard success and three ablation counterexamples, establishing corresponding guard requirements within the declared bounds.

Table 7. Bounded contract-lifecycle exploration and guard ablations.

Configuration	***States***	***Properties: outcome***
Full guards	*1,348*	*P1–P6: All hold; 2,320 transitions; no counterexample.*
No duplicate detection	*3,552*	*P2: 5-step violation; replay reactivates completed contract.*
No conflict precedence	*1,568*	*P5: 7-step violation; conflicting contracts active together.*
No forced expiration	*784*	*P3: 6-step violation; contract remains active past TTL.*

4.5. Deployment and limitations

Centralization eases global reasoning but adds latency/availability risk; distribution requires conflict/epoch management. Hybrid designs keep time-critical safety local. Estimator, selector, adapter, and monitor privileges are separated; outcomes are observed independently and operator attention is budgeted.

The lifecycle model omits estimator behaviour, delay distributions, and multi-domain authority. Mission weighting, correlated evidence, delayed effects, missing causal edges, and compromised targets threaten validity. Protocol heterogeneity demands recalibration; purposive selection and class coding limit generalization. Privacy, retention, emergency access, and role separation remain requirements. High resilience cannot override unrepresented safety/security violations.

5. CONCLUSION

Evidence qualification, causal dependencies, mission-quality measures, and action contracts make cross-layer coordination traceable. The theoretical contribution is the coordination model (1)–(12), the contract mechanism, and admissibility and stability criteria, and the results support three conclusions. First, because every record is weighted by (2)–(3) before estimation (4), replayed, stale, or inconsistent evidence loses influence by construction instead of steering the response. Second, the capped measure (7) lies in [0, 1] (Proposition 1), so coordinated and local policies can be compared through (11) under identical disturbances and resources. Third, within the declared bounds, duplicate detection, conflict precedence, and forced expiration are each necessary: P1–P6 held over 1,348 states and 2,320 transitions with all guards, whereas removing any one of them produced a counterexample. In terms of scope, the framework is the only entry in Table 5 that treats mission trajectory, contracts, and safety and stability explicitly together. The practical contribution is limited to application scenarios. Mission benefit, latency, overhead, and operational value remain to be established by simulation (V1), hardware-in-the-loop testing (V2), and pilot deployment (V3), followed by independent replication (V4–V5) with negative results reported.

ACKNOWLEDGEMENT

ChatGPT and Claude assisted only with English editing, reference formatting, and figure code; the authors verified all outputs and take full responsibility.

REFERENCES

[1] Sepúlveda Estay, D.A., Sahay, R., Barfod, M.B., Jensen, C.D. A systematic review of cyber-resilience assessment frameworks. *Computers & Security*, vol. 97, 2020, art. 101996. DOI: 10.1016/j.cose.2020.101996.

[2] Alhidaifi, S.M., Asghar, M.R., Ansari, I.S. A survey on cyber resilience: Key strategies, research challenges, and future directions. *ACM Computing Surveys*, vol. 56, no. 8, 2024, art. 196, pp. 1–48. DOI: 10.1145/3649218.

[3] Cho, H., Sung, J.-H., Kang, H.-J., Jang, J., Shin, D. Quantifying cyber resilience: A framework based on availability metrics and AUC-based normalization. *Electronics*, vol. 14, no. 12, 2025, art. 2465. DOI: 10.3390/electronics14122465.

[4] Sah, D.K., Amgoth, T. Parametric survey on cross-layer designs for wireless sensor networks. *Computer Science Review*, vol. 27, 2018, pp. 112–134. DOI: 10.1016/j.cosrev.2017.12.002.

[5] Stoynov, V. Explainable cross-layer machine learning for enterprise WLAN auditing. *International Journal on Information Technologies and Security*, vol. 18, no. 3, 2026, pp. 25–36. DOI: 10.59035/FYPL5639.

[6] AlSabeh, A., Khoury, J., Kfoury, E., Crichigno, J., Bou-Harb, E. A survey on security applications of P4 programmable switches and a STRIDE-based vulnerability assessment. *Computer Networks*, vol. 207, 2022, art. 108800. DOI: 10.1016/j.comnet.2022.108800.

[7] Hauser, F., Häberle, M., Merling, D., Lindner, S., Gurevich, V., Zeiger, F., Frank, R., Menth, M. A survey on data plane programming with P4: Fundamentals, advances, and applied research. *Journal of Network and Computer Applications*, vol. 212, 2023, art. 103561. DOI: 10.1016/j.jnca.2022.103561.

[8] Syed, N.F., Shah, S.W., Shaghaghi, A., Anwar, A., Baig, Z., Doss, R. Zero Trust Architecture: A comprehensive survey. *IEEE Access*, vol. 10, 2022, pp. 57143–57179. DOI: 10.1109/ACCESS.2022.3174679.

[9] Raza, S.M., Minerva, R., Crespi, N., Alvi, M., Herath, M., Dutta, H. A comprehensive survey of network digital twin architecture, capabilities, challenges, and requirements for edge–cloud continuum. *Computer Communications*, vol. 236, 2025, art. 108144. DOI: 10.1016/j.comcom.2025.108144.

[10] Verdecchia, R., Scommegna, L., Picano, B., Becattini, M., Vicario, E. Network digital twins: A systematic review. *IEEE Access*, vol. 12, 2024, pp. 145400–145416. DOI: 10.1109/ACCESS.2024.3453034.

[11] Scalingi, A., D'Oro, S., Restuccia, F., Melodia, T., Giustiniano, D. Det-RAN: Data-driven cross-layer real-time attack detection in 5G Open RANs. In: *Proceedings of the IEEE International Conference on Computer Communications (IEEE INFOCOM 2024)*, Vancouver, BC, Canada, 20–23 May 2024, pp. 41–50. DOI: 10.1109/INFOCOM52122.2024.10621223.

[12] Li, D.C., Tu, H.-H., Chou, L.-D. Cross-layer detection and defence mechanism against DDoS and DRDoS attacks in software-defined networks using P4 switches. *Computers and Electrical Engineering*, vol. 118, 2024, art. 109307. DOI: 10.1016/j.compeleceng.2024.109307.

[13] Yang, L., Naser, S., Shami, A., Muhaidat, S., Ong, L., Debbah, M. Toward zero touch networks: Cross-layer automated security solutions for 6G wireless networks. *IEEE Transactions on Communications*, vol. 73, no. 9, 2025, pp. 7650–7679. DOI: 10.1109/TCOMM.2025.3547764.

[14] Prokopovych-Tkachenko, D.I., Galushchenko, O., Kovalenko, Y., Mormul, M., Sarychev, V.I., Herasymenko, A. Spatial localization of Wi-Fi threat sources in critical infrastructure via RADIO-SOC event correlation. *International Journal on Information Technologies and Security*, vol. 18, no. 3, 2026, pp. 101–110. DOI: 10.59035/MKMU8096.

[15] Bocci, A., Forti, S., Ferrari, G.-L., Brogi, A. Declarative secure placement of FaaS orchestrations in the cloud-edge continuum. *Electronics*, vol. 12, no. 6, 2023, art. 1332. DOI: 10.3390/electronics12061332.

[16] Ross, R., Pillitteri, V., Graubart, R., Bodeau, D., McQuaid, R. Developing Cyber-Resilient Systems: A Systems Security Engineering Approach. NIST Special Publication 800-160, Volume 2, Revision 1. National Institute of Standards and Technology, Gaithersburg, MD, 2021. DOI: 10.6028/NIST.SP.800-160v2r1.

[17] Rose, S., Borchert, O., Mitchell, S., Connelly, S. Zero Trust Architecture. NIST Special Publication 800-207. National Institute of Standards and Technology, Gaithersburg, MD, 2020. DOI: 10.6028/NIST.SP.800-207.

[18] Pascoe, C., Quinn, S., Scarfone, K. The NIST Cybersecurity Framework (CSF) 2.0. NIST Cybersecurity White Paper 29. National Institute of Standards and Technology, Gaithersburg, MD, 2024. DOI: 10.6028/NIST.CSWP.29.

Information about the authors:

Oleksandr Korchenko — First Vice-Rector, State University of Information and Communication Technologies (SUICT), Kyiv, Ukraine. Doctor of Technical Sciences, Professor; Corresponding Member of the National Academy of Sciences of Ukraine; State Prize Laureate; Honored Worker of Science and Technology of Ukraine. ORCID: 0000-0003-3376-0631; e-mail: o.korchenko@duikt.edu.ua.

Dmytro Prokopovych-Tkachenko — PhD in Technical Sciences, Associate Professor; Head of the Department of Cybersecurity and Information Technologies, University of Customs and Finance, Dnipro, Ukraine; Doctoral Candidate, Department of Cybersecurity Systems and Technologies, SUICT, Kyiv, Ukraine. ORCID: 0000-0002-6590-3898; e-mail: omega2417@gmail.com.

Svitlana Kazmirchuk — Doctor of Technical Sciences, Professor; Professor of the Department of Cybersecurity Systems and Technologies, SUICT, Kyiv, Ukraine. Corresponding author. ORCID: 0000-0001-6083-251X; e-mail: s.kazmirchuk@duikt.edu.ua.

Mukafat Askerov — Candidate of Technical Sciences, Professor, Department of Mobile and Video Information Technologies, SUICT, Kyiv, Ukraine. ORCID: 0009-0000-6611-2732; e-mail: mukafat_ask@ukr.net.

Iurii Zhurov — Founder and Principal Consultant, SMB Cybersecurity Advisors LLC, Palm Harbor, Florida, USA. ORCID: 0000-0002-8276-2230; e-mail: iurii.zhurov@smbsecurityadvisors.com.

Manuscript received on 21 September 2026