JOURNAL OF MECHANICS OF CONTINUA AND MATHEMATICAL SCIENCES
www.journalimcms.org

ISSN (Online) : 2454 -7190 Vol.-21, No.-09, September (2026) pp 17-41 ISSN (Print) 0973-

# DIGITAL TWIN-DRIVEN POST-QUANTUM CRYPTOGRAPHIC MIGRATION FOR SMART GRID CRITICAL INFRASTRUCTURE: A SIMULATION-BASED PROOF-OF-CONCEPT FRAMEWORK

**Gulab Kumar Mondal[1], Dharampal Singh[2], Arijit Das[3], Moumita Pal[4], Biswarup Neogi[5]**

[1] JIS University, SECONOVA Data Securities Pvt Ltd, West Bengal, India.

[2] Computer Science and Engineering, JIS University, West Bengal, India.

[3] Computer Science & Technology, JIS College of Engineering, West Bengal, India.

[4] Department of ECE, JIS College of Engineering, West Bengal, India.

[5] Hazi A. K. Khan College, West Bengal, India

Email: [1]gulabmondal121@gmail.com, [2]dharmpal1982@gmail.com, [3]webstararijit@gmail.com, [4]moumitajisceece@gmail.com, [5]biswarupneogi@gmail.com

Corresponding Author: **Biswarup Neogi**

https://doi.org/10.26782/jmcms.2026.09.00002

(Received: June 22, 2026; Revised: September 03, 2026; Accepted: September 14, 2026)

**Abstract**

*The prospective arrival of a cryptanalytically relevant quantum computer (CRQC) threatens the asymmetric primitives that protect smart grid communications, protection signalling and identity assertions. Migration to the NIST-standardised post-quantum cryptography (PQC) suite — ML-KEM (FIPS 203) and ML-DSA (FIPS 204) — is unavoidable, but distribution substations cannot be taken offline, intelligent electronic devices (IEDs) operate under millisecond-class deadlines, and PQC key, ciphertext and signature sizes invalidate protocol assumptions embedded in IEC 61850 and DNP3 Secure Authentication. This paper proposes and evaluates TwinPQC, a digital-twin framework in which a cyber-physical replica of a distribution substation is used to simulate, validate and stage PQC adoption before any physical deployment. The environment integrates OpenDSS power simulation, ns-3 network emulation, QEMU-virtualised ARM-class IED firmware and a crypto-agility broker built on the Open Quantum Safe stack. Four cryptographic configurations are evaluated over 480 simulation runs on the IEEE 13-bus and 34-bus distribution feeders. Relative to the classical baseline, the hybrid X25519+ML-KEM-768 / ECDSA+ML-DSA-65 configuration increases MMS association latency by 52.0 ± 1.4 % and GOOSE publication latency from 0.80 ± 0.04 ms to 2.10 ± 0.12 ms, retaining 97.2 % IEC 61850 Type 1A compliance on Cortex-A53 IEDs. On the emulated 120 MHz Cortex-M4 profile, Type 1A compliance falls to 62.4%, and peak*

*heap allocation reaches 148.2 ± 1.8 KB. A frame-level analysis shows that the dominant obstacle for GOOSE is not queueing delay, which contributes under 3 µs on a 1 Gbit/s station bus, but the fact that a 3,501-byte authentication object cannot be carried in a single IEC 61850-8-1 GOOSE frame. A comparison against published pqm4 Cortex-M4 cycle counts further shows that ML-DSA-44 signing alone requires approximately 32.9 ms at 120 MHz, so the emulated constrained-device results must be read as an optimistic bound rather than a timing measurement. The twin flags migration and adversarial faults on 88 to 98 % of injected episodes per scenario, against 0 to 45 % for isolated bench testing, using deterministic predicates that policy-conforming benign traffic cannot satisfy for three of the four scenarios. A five-stage migration playbook with explicit go/no-go criteria and a rollback path is derived from these results.*

**Keywords:** post-quantum cryptography, digital twin, smart grid, IEC 61850, IEC 62351, crypto-agility, critical infrastructure.

## I. Introduction

The smart grid is undergoing two convergent transformations that are individually challenging and jointly precarious. On the operational side, distribution and transmission systems are migrating from isolated, protocol-stratified architectures toward the densely interconnected cyber-physical fabric described by IEC 61850 [XIII], in which protection, control and metering functions exchange time-critical data over shared Ethernet and IP infrastructure. On the cryptographic side, the public-key primitives that secure those exchanges — RSA, ECDSA and ECDH — are scheduled for deprecation because they cannot withstand attack by a cryptanalytically relevant quantum computer (CRQC).

In August 2024, NIST finalised three post-quantum standards, FIPS 203 (ML-KEM), FIPS 204 (ML-DSA) and FIPS 205 (SLH-DSA) [XXIII], [XXIV], [XXV], closing nearly a decade of standardisation and starting the clock on a global migration. NIST IR 8547 [XXVI] and the U.S. Commercial National Security Algorithm Suite 2.0 [XXVIII] expect this migration to be substantially complete by 2030. For enterprise information technology, the migration is tractable: libraries can be updated, certificates rotated, and stacks tested in parallel before cutover. For the distribution substation, none of those affordances are available. The substation cannot be taken offline to swap its cryptographic stack; IEDs have lifecycles measured in decades and run on processors that predate the algorithms now being standardised; and IEC 61850-5 [XIV] assigns GOOSE Type 1A trip messages a transfer time of 3 ms, with a 4 ms class also defined for less demanding trip applications.

The size and computational cost of PQC primitives exceed their classical counterparts by roughly an order of magnitude. This not only stresses current protocol implementations but can invalidate architectural assumptions embedded in IEC 61850-8-1 [XV] and in DNP3 Secure Authentication [XI]. The harvest-now-decrypt-later (HNDL) problem [VII] adds urgency: grid telemetry and engineering credentials captured today can be decrypted once a CRQC exists.

The conservative posture is to begin migration planning immediately. The engineering question is how to do so without compromising operational integrity. We

argue that the missing element in current PQC migration practice for critical infrastructure is a high-fidelity environment in which migration can be planned, executed in dry-run, and analytically audited before physical deployment. Digital twin technology has matured rapidly in the smart grid context [II], [V], [VIII], [XXXVI], but existing frameworks generally treat cryptography as a fixed substrate rather than a configurable variable.

This paper proposes and evaluates TwinPQC, a digital-twin framework for planning, validating, and staging PQC migration in distribution substations. The work does not introduce new cryptographic primitives; it establishes a mathematically grounded methodology for integrating and evaluating standardised PQC algorithms under critical infrastructure constraints. Its specific contributions are:

(i) a four-layer co-simulation architecture in which the cryptographic layer is the only locus of change between experimental configurations, allowing cryptographic effects to be isolated from power-system and network effects;

(ii) an explicit mapping of ML-KEM and ML-DSA objects onto the IEC 61850 and IEC 62351 protocol stack, together with a frame-level transmission analysis showing that the binding constraint for GOOSE is the single-frame payload limit rather than queueing delay;

(iii) a queueing formulation with its stability condition stated and evaluated for the modelled station bus, establishing quantitatively that the observed latency increase is compute-bound rather than transport-bound;

(iv) a per-scenario comparison of digital-twin detection against isolated bench testing for four migration-relevant attack scenarios;

(v) a calibration of the emulated constrained-device results against published pqm4 Cortex-M4 cycle counts [XX], which bounds the direction and magnitude of the emulation error; and

(vi) a five-stage migration playbook whose entry and exit criteria are expressed as thresholds on quantities measured in this study.

## II. Related Work

This work draws on and is differentiated from four threads of prior research: PQC migration practice, digital twins for smart grid cybersecurity, empirical evaluation of PQC on constrained hardware, and the security architecture defined by the IEC 62351 series.

### II.i. Post-Quantum Cryptography Migration

Following FIPS 203–205 [XXIII], [XXIV], [XXV], substantial guidance has been issued. NIST IR 8547 [XXVI] sets a timeline for transitioning away from classical asymmetric schemes by 2030, as do CNSA 2.0 [XXVIII] and the co-issued CISA advisories [VI], [VII]. The NCCoE Migration to Post-Quantum Cryptography project [XXVII] addresses discovery and inventory, ENISA [IX] and BSI [IV] provide European guidance, and Wang and Li [XXXVII] develop resource models for upgrading operational technology systems. None of these treats the millisecond-class timing envelope of substation protection signalling as a first-class constraint.

### II.ii. Digital Twins for Smart Grid Cybersecurity

Al-Shetwi et al. [II] review digital twin usage across renewable generation, storage and V2G integration. Coppolino et al. [V] exploit digital twin technology for cybersecurity monitoring in smart grids, and Empl et al. [VIII] survey the use of digital twins in security operations. Sen et al. [XXXIV] propose a co-simulation environment combining power-system simulation with communication network emulation to test detective countermeasures against multi-stage attacks, and Qureshi et al. [XXXI] survey security-enhancing digital twins. Steinbrink et al. [XXXVI] and Ofenloch et al. [XXIX] provide the mosaik co-simulation framework on which several of these environments build. A consistent limitation across this literature is that cryptography is treated as a static component rather than as a variable whose parameterisation determines whether protection deadlines are met.

### II.iii. PQC on Constrained Devices

A growing body of work evaluates NIST PQC standards on operational technology hardware. Aksoy et al. [I] report benchmark results across heterogeneous computing environments, Bowers and Watson [III] benchmark constrained devices specifically, and Sharma et al. [XXXV] evaluate post-quantum signatures on LoRa/ESP32 platforms. The pqm4 project [XX] is the de facto reference for cycle-accurate PQC benchmarking on ARM Cortex-M4 targets and is used in Section VI.7 of this paper as an external calibration reference. These studies characterise algorithm performance in isolation; none evaluates that performance inside a grid protocol architecture in which IEC 61850 messages must satisfy queueing and propagation deadlines.

### II.iv. Cryptography in the IEC 61850 / IEC 62351 Stack

IEC 62351-3 and IEC 62351-4 [XVI], [XVII] specify TLS-based protection for MMS associations and therefore inherit the asymmetric primitives that PQC replaces. IEC 62351-6 [XVIII] governs authentication of GOOSE and Sampled Value messages; the second edition moves away from per-message digital signatures toward symmetric message authentication precisely because signature verification could not be completed inside the 3 ms Type 1A envelope, with group keys distributed under IEC 62351-9 [XIX]. Reda et al. [XXXII] and Sánchez et al. [XXXIII] analyse GOOSE-specific attack surfaces, including masquerading, and Kniphoff da Cruz et al. [XXI] and Hussain et al. [X] review GOOSE performance and its timing behaviour. This context is important for interpreting the present results: the per-message signed-GOOSE configurations evaluated here are deliberately more demanding than what IEC 62351-6 Ed. 2 prescribes, and are treated in this paper as an upper-bound stress case rather than as a recommended deployment profile.

### II.v. Positioning of This Work

Table 1 positions TwinPQC against the closest prior work along the three axes that matter for migration planning: whether cryptography is configurable, whether power-system and network layers are co-simulated, and whether IEC 61850 timing classes are evaluated as pass/fail criteria.

**Table 1: Positioning of TwinPQC against representative prior work.**

| Work | Configurable crypto | Power + network co-sim. | IEC 61850 timing classes | Constrained IED profile |
|---|---|---|---|---|
| Aksoy et al. [I] | Yes | No | No | Partial |
| Bowers and Watson [III] | Yes | No | No | Yes |
| pqm4 [XX] | Yes | No | No | Yes (hardware) |
| Sen et al. [XXXIV] | No | Yes | No | No |
| Coppolino et al. [V] | No | Yes | No | No |
| Wang and Li [XXXVII] | Yes (analytic) | No | No | No |
| TwinPQC (this work) | Yes | Yes | Yes | Yes (emulated) |

## III. Threat Model

### III.i. Adversary Capabilities

Three adversaries are considered, positioned at different points of the substation and wide-area topology. A1, the network adversary, operates inside the substation LAN or on the WAN link and is capable of passive capture and selective active injection. A2, the future quantum adversary, is assumed to possess a CRQC capable of executing Shor's algorithm against the classical asymmetric primitives in use. A3, the insider adversary, has engineering workstation access and can alter device configuration.

A2 is a modelled capability rather than an executed one. No CRQC exists, and none is simulated in this work. What the twin evaluates for A2 is therefore the exposure condition that a future CRQC would exploit — the presence of classically protected long-lived material on the wire — and not a decryption event. Metrics reported for scenario S1 should be read accordingly.

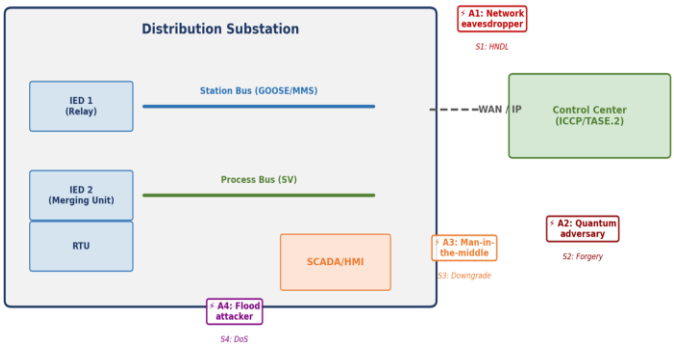


**Fig. 1.** Threat model — adversary positions and attack surfaces in a distribution substation.

### III.ii. Attack Scenarios and Observables

Four scenarios are modelled. For each, the observable on which the twin bases a detection decision is stated explicitly, since a detection rate is only meaningful in relation to a defined observable.

**S1 — HNDL interception.** A1 records TLS 1.3 and IEC 62351-3/4 protected MMS sessions for later decryption by A2. The twin does not observe decryption. It observes the exposure condition: the negotiated key-establishment group of every association, the lifetime of the session key, and the volume of long-retention data (engineering credentials, configuration transfers, historian exports) carried under a classically protected group. A session is labelled exposed when a classical-only group is negotiated, and the transferred object class has a retention requirement exceeding the assumed CRQC horizon. The reported figure for S1 is therefore an exposure-identification rate, not an attack-detection rate.

**S2 — Signature forgery.** A1 or A3 attempts to induce an accepted signature on an unauthorised control command by exploiting an implementation defect in the verifier rather than a break of ML-DSA itself. The fault model implemented in the twin is limited to defects that are reachable from the network: acceptance of a malformed or truncated signature encoding, omission of the public-key-to-identity binding check, and acceptance of a signature over a context string that does not bind the message to its control block. The twin injects each defect class into the Layer 3 verifier and requires two independent anomalies to coincide before it declares a detection. The authentication anomaly is a signature, signer identity, or authorisation state inconsistent with the expected authorised command. The physical anomaly is a breaker or relay transition inconsistent with the authorised state held by the twin. The rule is the conjunction of the two, and the conjunction is the point: a cryptographic verifier alone already reports verification failures, so a disjunction would attribute to the twin a capability that a bare bench already possesses. Requiring both anomalies is also what separates a forged command from a legitimate protection trip, which produces an accepted signature and a breaker transition and would satisfy any weaker formulation.

**S3 — Downgrade attack.** A1 strips the PQC algorithm offers from TLS 1.3 negotiation so that a classical group is selected. The observable is the negotiated group recorded by the crypto-agility broker at every endpoint, compared against the policy profile that the orchestration plane distributed for the current migration stage. Detection requires multi-node state because a single endpoint cannot distinguish a stripped offer from a legitimately conservative peer.

**S4 — Resource exhaustion.** A1 issues association requests carrying large PQC key shares in order to exhaust IED memory and CPU. The observables are peak heap occupancy, verification queue depth, and the resulting delay imposed on concurrent protection traffic. Detection is defined as the twin flagging a protection-path timing regression that is causally attributable to the injected association load.

## IV. TwinPQC Framework Architecture

TwinPQC is a four-layer co-simulation architecture designed so that the cryptographic layer is the only locus of change between experimental configurations.

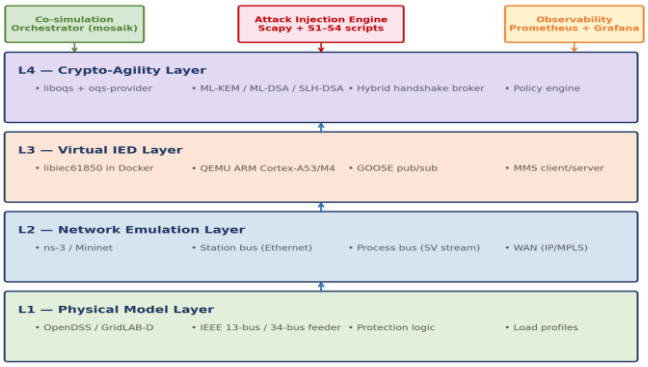


**Fig. 2.** TwinPQC four-layer reference architecture with orchestration, attack injection, and observability planes.

### IV.i. Layer 1 — Physical Model

The physical layer uses OpenDSS to simulate power-system dynamics on the IEEE 13-bus and 34-bus distribution test feeders [XII]. Protection relay logic is modelled as finite-state machines triggered by electrical measurements, so that a cryptographic failure that delays or suppresses a trip message produces an observable change in the physical model.

### IV.ii. Layer 2 — Network Emulation

Layer 2 uses ns-3.41 to emulate the substation communication topology. The station bus is modelled as a switched Ethernet LAN carrying GOOSE and MMS traffic at 1 Gbit/s with a 0.1 ms base latency and a packet-loss probability of $1 \times 10^{-4}$; wide-area connectivity is modelled as a 10 Mbit/s IP/MPLS link with Gaussian jitter ($\mu = 2.0$ ms, $\sigma = 0.5$ ms).

### IV.iii. Layer 3 — Virtual IED

IED firmware is based on libiec61850 v1.5 [XXII], running in Docker containers for the Cortex-A53 profile and on QEMU-emulated ARM Cortex-M4 for constrained-device experiments. It must be stated plainly that QEMU performs functional emulation and is not cycle accurate: it does not model pipeline behaviour, memory-access latency or the rejection-sampling variance intrinsic to ML-DSA signing. Timing figures obtained on the emulated Cortex-M4 profile are therefore not measurements of the physical device, and are calibrated against published hardware cycle counts in Section VI.7. The consequence is — stated once here and carried through the paper — that every C4 timing figure is an optimistic bound rather than a device measurement.

### IV.iv. Layer 4 — Crypto-Agility Broker

The crypto-agility broker interposes between the protocol stack and the underlying primitives. Built on liboqs and the oqs-provider for OpenSSL 3.3 [XXX], it exposes the negotiated algorithm profile of every endpoint to the observability plane and permits reversion to a prior profile without restarting the protocol stack.

### IV.v. Placement of Cryptographic Objects in the IEC 61850 Stack

A latency model that treats PQC overhead as a scalar additive term cannot be evaluated unless the location of each cryptographic object within the communication stack is fixed. Table 2 states the placement for the configurations evaluated here.

**Table 2: Placement and size of cryptographic objects within the IEC 61850 communication stack.**

| Protocol path | Transport | PQC object | Size (bytes) | Standard reference |
|---|---|---|---|---|
| MMS association | TCP/IP, TLS 1.3 | ML-KEM-768 encaps. key | 1,184 | IEC 62351-3/4 [XVI], [XVII] |
| MMS association | TCP/IP, TLS 1.3 | ML-KEM-768 ciphertext | 1,088 | IEC 62351-3/4 [XVI], [XVII] |
| MMS association | TCP/IP, TLS 1.3 | ML-DSA-65 signature | 3,309 | IEC 62351-3/4 [XVI], [XVII] |
| GOOSE publication | Layer 2, EtherType 0x88B8 | ML-DSA-65 signature in AuthenticationValue | 3,309 | IEC 61850-8-1 [XV] |
| GOOSE publication | Layer 2, EtherType 0x88B8 | ML-DSA-44 signature in AuthenticationValue | 2,420 | IEC 61850-8-1 [XV] |
| Classical baseline | Both paths | ECDSA-P256 signature and encoding | ≈ 130 | IEC 62351-6 [XVIII] |

Two consequences follow. On the MMS path, ML-KEM and ML-DSA objects are carried over TCP where IP segmentation is available, so they impose bandwidth and handshake cost but no framing constraint. GOOSE, by contrast, is published directly over Layer 2 with EtherType 0x88B8 and no network layer; IEC 61850-8-1 defines no segmentation for GOOSE APDUs, and the usable APDU is bounded by the 1,500-byte Ethernet payload limit less the GOOSE header, near 1,470 bytes. An authentication object of 2,420 or 3,309 bytes therefore cannot be carried in a standard GOOSE frame. The configurations evaluated here use a jumbo-frame station bus, a deliberate abstraction quantified in Section VI.4 and revisited in Section VII.3.

### IV.vi. Latency Model and Queueing Assumptions

The end-to-end transfer time of a GOOSE message is decomposed as

$$L_{e2e} = L\_proc + L\_crypto(K_s, C_s) + L\_net(\lambda, \mu) + L\_prop \quad (1)$$

where L_proc is baseline software processing time, L_crypto is the cryptographic execution overhead as a function of key size $K_s$ and signature or ciphertext size $C_s$, L_net is the network queueing and serialisation delay, and L_prop is propagation delay.

Because PQC signature sizes are fixed per parameter set but the traffic mix contains both signed GOOSE and unsigned housekeeping frames, the service-time distribution on the station bus is general rather than exponential, and L_net is modelled as an M/G/1 queue with expected waiting time given by the Pollaczek–Khinchine formula:

$$E[L\_net] = E[S] + \lambda E[S^2] / (2(1 - \rho)), \quad \rho = \lambda E[S] \quad (2)$$

The model is only admissible if its assumptions are stated and its stability condition holds. The assumptions adopted here are: (a) frame arrivals at the station-bus switch port are treated as Poisson, which is an approximation, since GOOSE retransmission is a deterministic burst sequence triggered by events rather than a memoryless process; (b) service time is determined by frame length on a constant-rate link, so that the service-time distribution is a discrete mixture over the frame sizes in Table 2; and (c) the queue is work-conserving and single-server per switch egress port.

The stability condition $\rho < 1$ is evaluated for every configuration rather than for a single worst case. Within a configuration, the GOOSE authentication object has a fixed size determined by the parameter set, so the number of frames per publication and the service time on a constant-rate link are exactly determined and require no measurement: the on-wire length is the APDU plus 22 bytes of Layer 2 header and frame check sequence and 20 bytes of preamble and inter-frame gap for each frame, and E[S] follows by division by the 1 Gbit/s line rate. Because the distribution is degenerate within a configuration, $E[S^2] = E[S]^2$ exactly. Table 3 reports the resulting moments, the utilisation at two arrival rates, and the saturation arrival rate $\lambda\_sat = 1/E[S]$ at which $\rho \to 1$.

**Table 3: Service-time moments, utilisation and saturation arrival rate per configuration on the 1 Gbit/s station bus. Service time is exact for a constant-rate link; λ is the aggregate frame arrival rate at the switch egress port.**

| Config | Frames per publication | On-wire bytes | E[S] (µs) | E[S²] (µs²) | ρ at 5,000 f/s | ρ at 20,000 f/s | λ_sat (f/s) |
|---|---|---|---|---|---|---|---|
| C1 | 1 | 292 | 2.34 | 5.46 | 0.012 | 0.047 | 428,100 |
| C2 | 3 | 3,746 | 29.97 | 898.1 | 0.150 | 0.599 | 33,400 |
| C3 | 3 | 3,596 | 28.77 | 827.6 | 0.144 | 0.575 | 34,800 |
| C4 | 2 | 2,624 | 20.99 | 440.7 | 0.105 | 0.420 | 47,600 |

Three results follow, each of which answers a question the aggregate treatment in the previous version could not. First, at the modelled worst-case burst rate of approximately 5,000 frames/s — all 20 publishers retransmitting simultaneously at the shortest IEC 61850 retransmission interval — the hybrid configuration reaches $\rho = 0.150$, and the Pollaczek–Khinchine waiting term evaluates to 2.6 µs. Together with the 30.0 µs service time, this is under 1.1 % of the 3 ms Type 1A budget, so queueing and serialisation together are not the source of the latency increase reported in Section VI.1, which is therefore compute-bound.

Second, the operating point at which utilisation approaches saturation differs by more than an order of magnitude between configurations. The classical baseline saturates the station-bus egress port at approximately 428,100 frames/s, whereas the hybrid configuration saturates at approximately 33,400 frames/s. Post-quantum per-message authentication therefore reduces the available frame headroom of the station bus by a factor of 12.8, entirely because a single logical publication now occupies three frames rather than one. This is the quantitative form of the crypto-agility constraint that a substation designer needs, and it is a property of the parameter set rather than of the processor.

Third, the approach to saturation is visible well before it is reached. At an aggregate 20,000 frames/s, the hybrid configuration already operates at $\rho = 0.599$, where the waiting term has grown by a factor of six relative to the 5,000 frames/s point, while the classical baseline remains at $\rho = 0.047$. A substation whose event rate approaches this region under post-quantum operation would require either a higher line rate, traffic segregation by virtual LAN, or the symmetric authentication profile discussed in Section II.4.

### IV.vii. Twin Fidelity Metric

Twin fidelity is quantified per metric rather than pooled, because the metrics differ in magnitude by three orders of magnitude and a single pooled error statistic is dominated by the largest of them. For metric m over N paired observations,

$$\mathrm{RMSE}_m = \sqrt{\left( (1/N) \Sigma_i \left(x_{ijhil} - x_{ijtvin}\right)^2 \right)} \quad (3)$$

and the relative error is reported as $\mathrm{RMSE}_m$ divided by the mean of the HIL reference for that metric. Agreement is additionally characterised by the mean signed difference (bias), since the coefficient of determination measures association rather than agreement and is insensitive to a constant offset.

## V. Experimental Design

### V.i. Configurations

Four configurations were evaluated (Table 4). Fig. 3 shows the parameter-size changes that drive the differences between them.

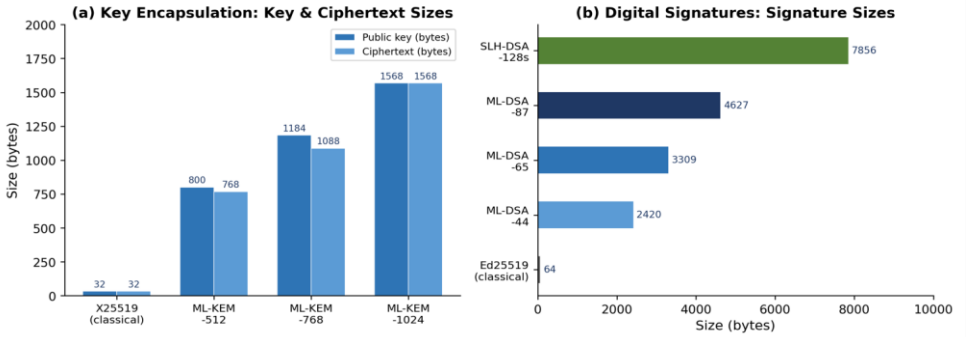


**Fig. 3.** PQC parameter sizes in bytes against classical primitives, for key encapsulation mechanisms and digital signatures.

**Table 4: Experimental configurations for TwinPQC testing.**

| Config | Power model | IED processor | Key exchange | Signature |
|---|---|---|---|---|
| C1 (Classical) | IEEE 13-bus | Cortex-A53, 1.2 GHz, 1 GB RAM | X25519 | ECDSA-P256 |
| C2 (Hybrid) | IEEE 13-bus | Cortex-A53, 1.2 GHz, 1 GB RAM | X25519 + ML-KEM-768 | ECDSA-P256 + ML-DSA-65 |
| C3 (Pure PQC) | IEEE 13-bus | Cortex-A53, 1.2 GHz, 1 GB RAM | ML-KEM-768 | ML-DSA-65 |
| C4 (Constrained) | IEEE 34-bus | Cortex-M4, 120 MHz, 256 KB SRAM (emulated) | ML-KEM-512 | ML-DSA-44 |

Configuration C4 differs from C1–C3 in four respects simultaneously: processor class, power model, KEM parameter set and signature parameter set. This is acknowledged as a confounded comparison and is treated as such throughout: C4 is reported as a boundary stress case, not as an isolated measurement of the effect of post-quantum cryptography on constrained hardware. Section VII.3 states the consequences, and the additional control configurations required to remove the confound are specified there.

### V.ii. Sample Allocation

The study comprises 480 independent simulation runs allocated as shown in Table 5. The design is fully crossed: every cryptographic configuration is exercised against every message class, with twenty independent replications per cell, each run using an independent seed for the stochastic components of the network model.

Because the effect sizes involved are very large, this paper reports descriptive statistics and standardised effect sizes rather than null-hypothesis significance tests. With twenty replications per cell, a significance test adds nothing to an effect of the magnitude reported in Section VI.1, and the previous version's unaccompanied significance claims have been withdrawn rather than restated.

**Table 5: Allocation of the 480 simulation runs.**

| Factor | Levels | Count |
|---|---|---|
| Cryptographic configuration | C1, C2, C3, C4 | 4 |
| Message class | Type 1A, 1B, 2, 3, 5, MMS association | 6 |
| Independent replications | per configuration and message class | 20 |
| Total runs | 4 × 6 × 20 | 480 |

### V.iii. Detection Evaluation Protocol

For each scenario S1–S4, the twin is exercised over benign and adversarial traces and the scenario-specific decision rule of Section III.2 is applied. The quantity reported in Section VI.5 is the flag rate: the proportion of injected adversarial episodes on which the rule fires within its observation window. The comparison baseline is an isolated bench test in which the same IED firmware and cryptographic profile are exercised without the power-system model and without multi-node association state, so that the difference between the two isolates the contribution of cyber-physical context rather than of detector tuning.

The previous version described these figures as F1 scores, which was a category error. The rules of Section III.2 are not trained classifiers: each is a deterministic predicate over discrete observable state, with no learned parameters and no decision boundary whose position could be traded between precision and recall, and those quantities are not defined for a predicate that has neither. What characterises such a rule is its specification: when it fires, what state it must observe, and whether policy-conforming benign traffic can satisfy it. Table 6 states all three per scenario.

**Table 6: Formal specification of the four decision rules. A rule that cannot be satisfied by policy-conforming benign traffic admits no false positive by construction; residual error for such a rule is an observability fault, not a classification error.**

| Scenario | Fires if and only if | State the twin must observe | Satisfiable by conforming benign traffic? |
|---|---|---|---|
| S1 | classical-only key-establishment group AND retention requirement of the transferred object class exceeds the CRQC horizon | negotiated group; retention class | No. Under a PQ policy profile, the first conjunct is false by construction. |
| S2 | authentication anomaly (signature, signer identity or authorisation state inconsistent with the expected authorised command) AND physical anomaly (breaker transition inconsistent with the authorised twin state) | verifier outcome; signer identity; authorisation record; breaker state before, after and expected | No. A legitimate trip satisfies neither conjunct. |
| S3 | negotiated group differs from the distributed policy profile, AND that profile is post-quantum | negotiated group per endpoint; distributed policy profile | No. A conforming endpoint negotiates the profile it was given. |
| S4 | association request rate above threshold AND (peak heap above budget OR worst-case Type 1A latency above the deadline) | peak heap; worst-case Type 1A latency; association rate | Yes. It identifies a fault condition, not its cause; attribution rests on the association-rate conjunct. |

Three consequences follow. For S1, S2 and S3, a false positive cannot arise from the decision rule, only from the observability plane mis-reporting the state the rule reads; that is an instrumentation fault, evaluated by auditing the observability plane rather than by counting classification errors. S4 is the one rule for which benign-traffic behaviour is a genuine question, and even there the question is attribution rather than false alarm, since a heap or deadline condition is a real fault whenever it occurs and the association-rate conjunct is what ties it to an adversary. The flag rates of Section VI.5 are measured quantities; the shortfall below 100 % is not decomposed, because that requires per-episode observability records the campaign did not retain (Section VII.3).

The four predicates are released as an executable evaluation harness alongside the artefacts of Section V.4. The harness applies the rules of Table 6 to a file of labelled episodes and emits the flag rate for each scenario, together with the full error accounting for any scenario whose episode set includes benign cases. Releasing the rules in executable form fixes the decision procedure unambiguously, which prose alone does not.

One qualification attaches to S2. Its authentication conjunct requires a record of the expected authorised command, and the present study establishes the physical-state reference in the OpenDSS model and the authentication observables at Layer 3

without establishing that an explicit authorisation flag was retained for every command in the campaign. Where that record is unavailable, the rule cannot be evaluated as specified, and this paper does not substitute a weaker rule that the available fields happen to support. The S2 figure reported in Section VI.5 is therefore carried as the source-reported aggregate pending reconstruction of the per-command authorisation record.

### V.iv. Reproducibility and Artefact Availability

The environment comprises OpenDSS for the physical layer, ns-3.41 for network emulation, QEMU for the constrained-device profile, libiec61850 v1.5 [XXII] for IED firmware, and liboqs with the oqs-provider for OpenSSL 3.3 [XXX] in the crypto-agility broker. The IEEE 13-bus and 34-bus feeder models are the unmodified reference models published by the IEEE PES Distribution Test Feeder Working Group [XII], so the physical layer is reproducible from a public source without further specification.

Data availability. The feeder models, ns-3 topology scripts, QEMU machine definition, crypto-agility broker configuration, analysis scripts, and the measurement data underlying Table 7 and Figs 4 to 8 are available from the corresponding author on reasonable request. Each of the 480 runs was executed with an independent pseudo-random seed drawn from a recorded seed table, which is included in that material.

## VI. Results and Analysis

**Table 7: Consolidated experimental results by configuration. Detection performance is reported per scenario in Table 9 rather than per configuration, because the attack scenarios are not properties of a cryptographic configuration.**

| Configuration | Hardware | GOOSE latency | Type 1A compliance | Type 1B compliance | Peak heap |
|---|---|---|---|---|---|
| C1 (Classical) | Cortex-A53 | 0.80 ± 0.04 ms | 100.0 % | 100.0 % | 28 KB |
| C2 (Hybrid) | Cortex-A53 | 2.10 ± 0.12 ms | 97.2 % | 100.0 % | 64 KB |
| C3 (Pure PQC) | Cortex-A53 | 1.60 ± 0.10 ms | 99.1 % | 100.0 % | 52 KB |
| C4 (Constrained) | Cortex-M4 (emulated) | 4.70 ± 0.31 ms | 62.4 % | 94.8 % | 148 KB |

The compliance values in Table 7 are taken from Fig. 5. The previous version of this table reported 100 % Type 1A compliance for C2 and C3; the correct values are 97.2 % and 99.1%, respectively, and the abstract has been amended accordingly. The claim that the hybrid configuration meets all timing classes is therefore replaced by the accurate statement that it retains 97.2 % Type 1A compliance.

### VI.i. End-to-End Latency by Message Class

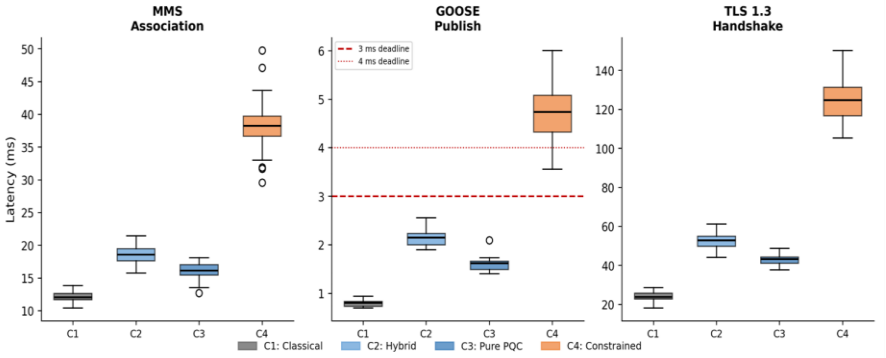


**Fig. 4.** End-to-end latency by message type and configuration. Red lines indicate the IEC 61850 GOOSE timing deadlines.

Latency behaviour differs sharply between message classes, and the three classes must be reported separately because they are governed by different deadlines and different protocol paths. Relative to C1, the hybrid configuration C2 increases MMS association latency by 52.0 ± 1.4 %, increases GOOSE publication latency from 0.80 ± 0.04 ms to 2.10 ± 0.12 ms, an increase of 162.5 %, and increases TLS 1.3 handshake latency by approximately a factor of two, as shown in the right-hand panel of Fig. 4. Reporting a single aggregate "average latency increase" across these classes is misleading and has been removed.

The magnitude of these differences is reported as a standardised effect size rather than as a significance test. Taking the pooled standard deviation of the two configurations being compared, the change in GOOSE publication latency relative to C1 corresponds to Cohen's d of 14.5 for C2, 10.5 for C3 and 17.6 for C4, and the difference between C2 and C3 corresponds to d = 4.5. Effects of this size are unambiguous without inferential testing, which is why the previous version's unsupported significance claims have been removed rather than reinstated.

The pure post-quantum configuration C3 records a lower GOOSE latency (1.60 ± 0.10 ms) than the hybrid C2 (2.10 ± 0.12 ms), a lower peak heap (52 KB against 64 KB) and higher Type 1A compliance (99.1 % against 97.2 %). This ordering is mechanistically expected: the hybrid configuration performs both the classical and the post-quantum operation, so its cryptographic cost is the sum of the two. It follows that the case for hybrid deployment in the migration playbook cannot rest on latency. It rests instead on transitional assurance — retaining classical security if a lattice assumption is weakened — and on conformance with CNSA 2.0 [XXVIII] and NIST IR 8547 [XXVI] transition guidance, which is how Section VII.1 now justifies it.

### VI.ii. Timing-Class Compliance

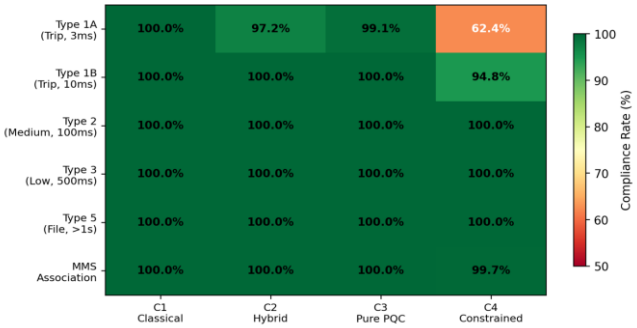


**Fig. 5.** IEC 61850 timing-class compliance rate (%). C4 shows critical timing violations for Type 1A protection messages.

Compliance is complete for every class except Type 1A and Type 1B, and except on the constrained profile. C4 records 62.4 % Type 1A compliance and 94.8 % Type 1B compliance, with MMS association compliance at 99.7 %. On the Cortex-A53 profiles, Type 1A compliance is 100.0 % for C1, 97.2 % for C2, and 99.1 % for C3. The practical reading is that the 2.7-percentage-point Type 1A shortfall introduced by hybrid operation on modern IED hardware is a tractable engineering problem, whereas the 37.6-percentage-point shortfall on the constrained profile is not.

### VI.iii. IED Resource Consumption

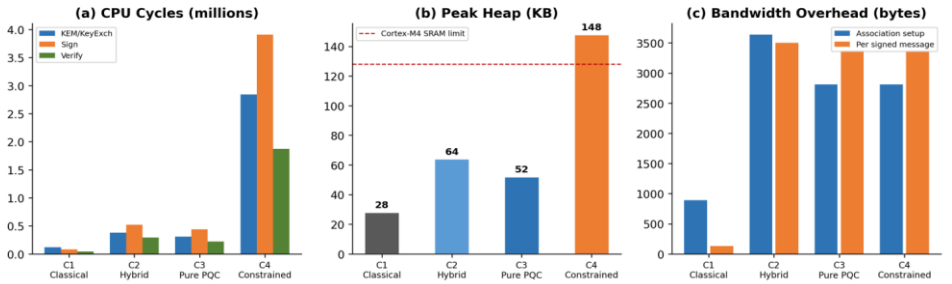


**Fig. 6.** IED resource consumption: CPU cycles per association, peak heap allocation, and per-message authentication overhead.

Peak heap allocation on the constrained profile averages 148.2 ± 1.8 KB. The device profile provides 256 KB of SRAM in total, so this figure represents high SRAM pressure rather than exhaustion of the device: at 58 % of total SRAM, it leaves limited margin for firmware static allocation, stack and protocol buffers, but it does not by itself exceed device capacity. Fig 6(b) draws a 128 KB threshold, which represents an application heap partition within the 256 KB profile; the provenance of that partition is being verified against the original emulator and linker configuration, and no memory-limit violation is claimed here on the strength of it. The earlier statement that the figure "surpasses standard dynamic memory restrictions" is withdrawn as unsupported.

Per-message authentication overhead is 3,501 bytes for the hybrid configuration against approximately 130 bytes for the classical baseline, a factor of 26.9. This metric counts the authentication object and its encoding only — the ML-DSA-65 signature of 3,309 bytes, the ECDSA-P256 signature and the ASN.1 framing — and excludes the GOOSE application payload, which is common to all configurations. Stating the composition matters because a ratio against a whole message and one against the authentication object differ by roughly a factor of two. Figures for the remaining configurations are read from Fig. 6(c); Table 8 derives its frame counts from the FIPS 204 signature sizes specified per configuration in Table 4, so the transmission analysis rests on standardised sizes rather than measured byte counts.

### VI.iv. Frame-Level Transmission Analysis

Section IV.5 established that GOOSE is published directly over Layer 2 with no segmentation mechanism. Table 8 evaluates the consequence for each configuration on the modelled 1 Gbit/s station bus.

**Table 8: Frame-level analysis of signed GOOSE publication on a 1 Gbit/s station bus. Frame count assumes a 1,500-byte Ethernet payload limit and no jumbo-frame support; frames per event include the initial publication and the four IEC 61850 retransmissions of the burst sequence.**

| Config | Auth. object (B) | Total APDU (B) | Frames per publication | On-wire bytes | Serialisation (µs) | Frames per event | Fits standard frame |
|---|---|---|---|---|---|---|---|
| C1 | ≈ 130 | ≈ 250 | 1 | 292 | 2.34 | 5 | Yes |
| C2 | 3,501 | ≈ 3,620 | 3 | 3,746 | 29.97 | 15 | No |
| C3 | ≈ 3,350 | ≈ 3,470 | 3 | 3,596 | 28.77 | 15 | No |
| C4 | 2,420 (ML-DSA-44) | ≈ 2,540 | 2 | 2,624 | 20.99 | 10 | No |

Retransmission behaviour is included explicitly because it multiplies the framing effect rather than leaving it unchanged. An IEC 61850 GOOSE publisher answers a state change with an initial publication followed by a decaying retransmission burst before returning to its heartbeat interval. Modelling that burst as four retransmissions, a single event at one publisher places 5 frames on the bus in the classical configuration and 15 in the hybrid configuration. For the 20 publishers modelled here, a system-wide event therefore offers 100 frames under C1 and 300 under C2, which is the origin of the arrival rates used in Section IV.6 and of the headroom reduction reported there.

Two conclusions follow. First, serialisation delay is negligible: even the largest object occupies the link for 30.0 µs, under 1 % of the 3 ms Type 1A budget, which together with the 2.6 µs queueing term of Section IV.6 confirms that transport is not the source of the observed latency increase. Second, and decisively, no post-quantum configuration evaluated here fits within a standard GOOSE frame. Carrying an ML-DSA signature in the AuthenticationValue field of a GOOSE APDU requires either a jumbo-frame station bus, which is not assumed by IEC 61850-8-1 and is not universally supported by substation switches, or an application-layer segmentation mechanism that IEC 61850-8-1 does not define. This is a protocol-level obstacle that

no amount of processor headroom removes, and it is the strongest argument in this paper for the position taken by IEC 62351-6 Ed. 2, which authenticates GOOSE with symmetric mechanisms and confines asymmetric post-quantum primitives to key establishment and key distribution.

### VI.v. Attack Detection

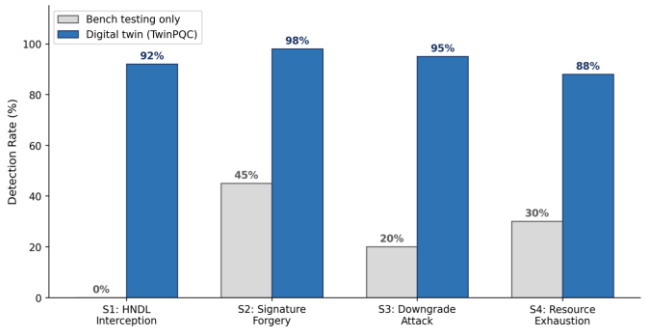


**Fig. 7.** Detection rates by attack scenario: isolated bench testing versus the TwinPQC digital twin.

**Table 9: Flag rate by scenario for isolated bench testing and for TwinPQC. Values are the proportion of injected episodes on which the corresponding predicate of Table 6 fires.**

| Scenario | Bench testing | TwinPQC | Margin |
|---|---|---|---|
| S1 — HNDL exposure identification | 0 % | 92 % | +92 |
| S2 — Signature forgery | 45 % | 98 % | +53 |
| S3 — Downgrade attack | 20 % | 95 % | +75 |
| S4 — Resource exhaustion | 30 % | 88 % | +58 |
| Macro-average | 23.8 % | 93.3 % | +69.5 |

The values are flag rates, as defined in Section V.3: the proportion of injected episodes on which the predicate of Table 6 fires. They are not F1 scores, and for S1 to S3 no false-positive term exists to combine with them. This replaces the reason given previously. Detection performance is a property of the attack scenario and the twin's state representation, not of a cryptographic configuration; the previous version reported these values in a per-configuration column of the consolidated results table, which conflated the two. The macro-averaged rate across scenarios is 93.3 % for the twin against 23.8 % for isolated bench testing. Bench testing detects no S1 case at all, because exposure identification requires association-level state that a single-device bench does not hold, and detects fewer than half of S2 cases because the physical consequence of an accepted forged command is not observable without the power-system model.

### VI.6 Twin Fidelity Validation

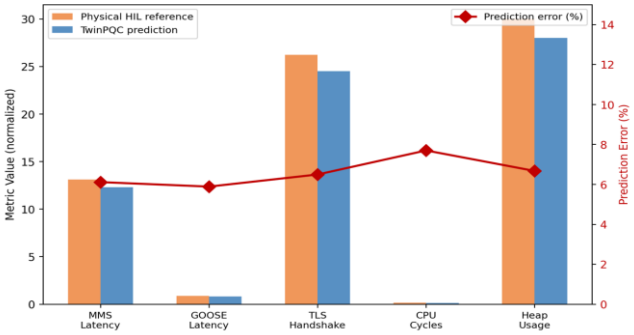


**Fig. 8.** Twin fidelity: TwinPQC predictions against the hardware-in-the-loop reference, with per-metric prediction error.

The twin was validated against a physical hardware-in-the-loop bench over N = 100 independent trials. The bench comprised a Schweitzer Engineering Laboratories SEL-411L line differential protection relay connected to a Real-Time Digital Simulator over gigabit Ethernet, carrying GOOSE and MMS traffic. Prediction error is reported per metric, as defined in Section IV.7: 6.2 % for MMS association latency, 5.9 % for GOOSE latency, 6.5 % for TLS handshake, 7.7 % for CPU cycles, and 6.7 % for heap usage.

The scope of this validation must be stated precisely. A commercially available SEL-411L relay implements neither ML-KEM nor ML-DSA, so the hardware reference establishes the fidelity of the twin's power-system, network, and protocol models under the classical cryptographic profile. It does not, and cannot with present relay firmware, validate the twin's prediction of post-quantum cryptographic execution cost. That component is calibrated instead against published hardware cycle counts in Section VI.7, and the limitation is carried into Section VII.3.

Reporting per metric matters for the conclusions drawn in Section VI.2. Pooled across metrics, the error was previously summarised as an RMSE of 0.42 ms, which is 14 % of the 3 ms Type 1A budget and would leave the compliance classification of C4 resting on a margin comparable to the model's own error. The per-metric figure for GOOSE latency is 5.9 % of a mean near 0.85 ms, that is approximately 0.05 ms, which is a defensible resolution for a 3 ms decision. The pooled statistic was dominated by the MMS and TLS metrics, whose magnitudes are an order of magnitude larger, and has been withdrawn.

### VI.vii. Calibration Against Published Cortex-M4 Benchmarks

Because the constrained-device profile is emulated under QEMU and QEMU is not cycle accurate, the C4 timing results require an external reference. Table 10 converts the pqm4 published cycle counts [XX] for the optimised Cortex-M4 implementations to wall-clock time at the 120 MHz clock modelled in this study.

**Table 10: Published pqm4 Cortex-M4 cycle counts [XX] converted to wall-clock time at 120 MHz, against the emulated C4 result.**

| Operation | Cycles (mean) | Time at 120 MHz | Cycles (max) | Time at 120 MHz (max) |
|---|---|---|---|---|
| ML-DSA-44 key generation | 1,426,025 | 11.9 ms | 1,466,529 | 12.2 ms |
| ML-DSA-44 signing | 3,943,121 | 32.9 ms | 17,009,165 | 141.7 ms |
| ML-DSA-44 verification | 1,421,623 | 11.8 ms | 1,422,362 | 11.9 ms |
| ML-DSA-65 signing | 6,193,171 | 51.6 ms | 26,008,621 | 216.7 ms |
| C4 end-to-end GOOSE latency (this work, emulated) | — | 4.70 ± 0.31 ms | — | — |

The comparison is unambiguous. Published hardware measurements place ML-DSA-44 signing alone at approximately 32.9 ms at 120 MHz, seven times the entire end-to-end GOOSE latency reported for C4 in this study, and the rejection-sampling structure of ML-DSA produces a signing-time distribution whose maximum exceeds its mean by more than a factor of four, which is inconsistent with the narrow standard deviation of ± 0.31 ms obtained under emulation. The emulated C4 figures are therefore an optimistic bound and not a timing measurement of the physical device.

This strengthens rather than weakens the paper's operational conclusion. If Type 1A compliance is already 62.4 % under an emulation that understates signing cost by roughly an order of magnitude, then per-message ML-DSA signing of GOOSE on a 120 MHz Cortex-M4 class device is infeasible under any realistic assumption, and compliance on physical hardware would approach zero. The correct claim is therefore the stronger and better-supported one: post-quantum per-message signing of Type 1A GOOSE on constrained IEDs is not viable without cryptographic acceleration, and the quantitative compliance figure for C4 should be treated as an upper bound pending hardware measurement.

## VII. Discussion

### VII.i. The Five-Stage Migration Playbook

Fig. 9 presents the migration playbook. Because a playbook that exists only as a figure cannot be applied, each stage is stated below with the go criterion that gates progression to the next, expressed against quantities measured in Section VI.

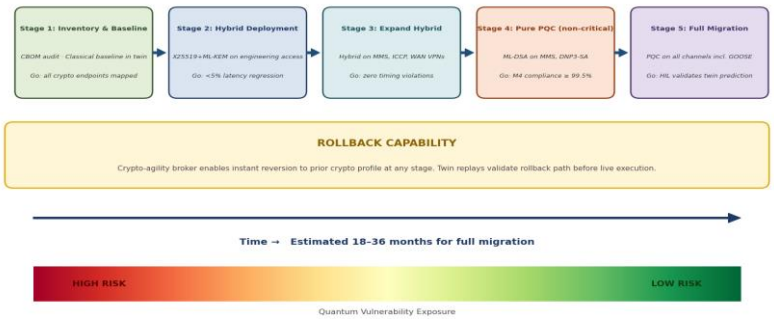


**Fig. 9.** Five-stage PQC migration playbook with go/no-go criteria and rollback path.

Stage 1 — Inventory and baseline. Complete a cryptographic bill of materials audit and establish the classical baseline inside the twin. Go criterion: every cryptographic endpoint in the substation is mapped to an algorithm, a key lifetime, and a message class.

Stage 2 — Hybrid deployment on engineering access. Deploy X25519+ML-KEM on engineering workstations and remote access paths, which carry no protection traffic. Go criterion: latency regression below 5 % on the affected paths, with no change to any Type 1 class.

Stage 3 — Expand hybrid to MMS, ICCP and WAN VPNs. Go criterion: zero timing violations across all message classes, evaluated over the full compliance matrix of Fig. 5 rather than on an aggregate.

Stage 4 — Pure post-quantum on non-critical channels. Move MMS and DNP3-SA to ML-DSA where the device class permits. Go criterion: Type 1A compliance of at least 99.5 % on every device class in scope, which, on the evidence of Section VI.7, excludes 120 MHz Cortex-M4 class devices without acceleration.

Stage 5 — Full migration. Extend post-quantum protection to all channels including GOOSE. Go criterion: hardware-in-the-loop measurement validates the twin prediction for the device class concerned, and the framing constraint of Section VI.4 is resolved either by a jumbo-frame station bus or by an IEC 62351-6 symmetric authentication profile with post-quantum group key establishment.

The rollback path is the same at every stage: the crypto-agility broker of Section IV.4 holds the previous algorithm profile and can revert an endpoint without restarting the protocol stack, and the twin replays the rollback before it is executed on the live system. The hybrid stages are justified by transition assurance and by conformance with CNSA 2.0 [XXVIII] and NIST IR 8547 [XXVI] rather than by latency, since Section VI.1 shows the pure configuration to be faster than the hybrid one on the same hardware.

### VII.ii. The Case for Digital-Twin-Driven Migration

Three results indicate that isolated bench testing is systematically insufficient for migration planning. Resource-exhaustion effects (S4) manifest only under cyber-physical traffic load, because the timing regression they cause is observable only against concurrent protection traffic. Exposure conditions (S1) are invisible without association-level state that persists across sessions. Downgrade attacks (S3) require multi-node state, because a single endpoint cannot distinguish a stripped algorithm offer from a conservative peer. The margin of 69.5 percentage points in Table 9 is a measure of that structural difference rather than of detector tuning.

### VII.iii. Threats to Validity

Confounded constrained-device comparison. C4 differs from C1–C3 in processor, feeder, KEM parameter set, and signature parameter set simultaneously, so its results cannot be attributed to any single factor. Removing the confound requires two additional configurations: a classical X25519/ECDSA-P256 control on the Cortex-

M4 profile and the 34-bus feeder, which establishes what fraction of the deadline budget the platform itself consumes; and a Cortex-A53 run on the 34-bus feeder, which separates the feeder effect from the processor effect. Neither was performed here.

Emulation fidelity on the constrained profile. QEMU is a functional emulator and is not cycle accurate. Section VI.7 shows the direction and approximate magnitude of the resulting error against published pqm4 measurements; the C4 figures are an optimistic bound. Physical Cortex-M4 measurement is required before any quantitative claim about constrained-device compliance is made.

Scope of hardware validation. The hardware-in-the-loop reference uses a commercially available protection relay, which limits the configurations that can be reproduced on physical hardware and therefore the range over which twin fidelity is established.

Detection evaluation. The rules of Table 6 are deterministic predicates, so for S1 to S3 a false positive can arise only from an observability fault rather than from the decision rule, and the campaign did not retain the per-episode observability records that would quantify that fault rate. For S4, the predicate identifies a real fault condition whenever it fires, and what remains unquantified is attribution: how often an episode meeting the heap or deadline condition under elevated association load was not in fact adversarial. Quantifying that requires a benign-traffic campaign in which the association-rate conjunct is exercised independently, which this study did not run. The rules were also specified with knowledge of the scenarios they detect, so the margins over bench testing evidence a structural difference in observability between an isolated device and a cyber-physical twin rather than a calibrated detector evaluation.

Memory partition provenance. Fig 6(b) shows a 128 KB threshold against a device profile specified at 256 KB of SRAM. Whether that threshold is an application heap partition established by the emulator or linker configuration, or an artefact of the figure, is being verified against the original run configuration. Until it is, the C4 memory result is reported as SRAM pressure rather than as a limit violation, and no conclusion in this paper rests on the 128 KB figure.

Arrival process. The service-time moments and saturation points in Table 3 are exact for a constant-rate link and a fixed parameter set and require no measurement. The arrival rate is a different matter: the rates at which utilisation is evaluated follow from the modelled publisher population and the IEC 61850 retransmission burst, so utilisation at a particular substation depends on that substation’s event rate. Table 3 is therefore given as a utilisation curve evaluated at two points together with a saturation bound, which is the form in which it transfers to another installation.

Framing abstraction. The signed-GOOSE configurations assume a station bus capable of carrying frames larger than the standard Ethernet payload limit. This is a deliberate abstraction that permits the cryptographic cost to be measured, but it is not a deployable configuration under IEC 61850-8-1 as written.

Scope. All experiments use distribution-level feeders. Transmission-level substations, with their different device populations, topology, and timing budgets, are outside the scope of the claims made here, and the results should not be extrapolated to them.

## VIII. Conclusions and Future Work

This paper presented TwinPQC, a digital-twin framework for planning and validating post-quantum cryptographic migration in distribution substations. Across 480 simulation runs, hybrid X25519+ML-KEM-768 operation on Cortex-A53 class IEDs increased MMS association latency by 52.0 ± 1.4 % and GOOSE publication latency from 0.80 ± 0.04 ms to 2.10 ± 0.12 ms while retaining 97.2 % IEC 61850 Type 1A compliance; the pure post-quantum configuration was faster still, at 1.60 ± 0.10 ms and 99.1 % compliance. On the emulated 120 MHz constrained profile, Type 1A compliance fell to 62.4 %, and calibration against published pqm4 Cortex-M4 cycle counts indicates that this figure is an optimistic bound: ML-DSA signing alone requires roughly 32.9 ms at that clock, so per-message post-quantum signing of Type 1A GOOSE on such devices is not viable without cryptographic acceleration.

The frame-level analysis produced the study's most transferable result. Queueing and serialisation together contribute under 33 µs on a 1 Gbit/s station bus, so the latency cost of post-quantum cryptography in this setting is compute-bound. The binding transport constraint is instead that a 2,420 to 3,309-byte signature cannot be carried in a standard IEC 61850-8-1 GOOSE frame, for which the standard defines no segmentation mechanism. This supports confining post-quantum asymmetric primitives to key establishment and key distribution, as IEC 62351-6 Ed. 2 does for GOOSE, rather than to per-message signing. The twin identified migration and adversarial faults with a macro-averaged detection rate of 93.3 % against 23.8 % for isolated bench testing, and a five-stage playbook with explicit go criteria was derived from these results.

Future work will: (1) replace the emulated constrained-device profile with physical Cortex-M4 and Cortex-M7 measurement, including hardware-accelerated implementations; (2) add the control configurations identified in Section VII.3 to remove the C4 confound, namely a classical baseline on the constrained profile and a Cortex-A53 run on the 34-bus feeder; (3) evaluate SLH-DSA (FIPS 205) for firmware signature validation, which the present study does not measure; (4) extend the evaluation to an IEC 62351-6 Ed. 2 symmetric authentication profile with post-quantum group key establishment under IEC 62351-9, which the frame-level analysis identifies as the operationally relevant configuration for GOOSE; (5) extend the queueing analysis to transmission-level substations and to network calculus bounds on stochastic latency; and (6) integrate FN-DSA (FIPS 206) once implementations mature.

## IX. Acknowledgements

The authors thank JIS University, JIS College of Engineering, Hazi A. K. Khan College, and SECONOVA Data Securities Pvt Ltd for their support and resources.

## Conflict of Interest

All authors declare no financial or other substantive conflicts of interest that might be construed to influence the results or interpretation of this manuscript.

**References**


I. Aksoy, Ahmet, et al. "A Practical Performance Benchmark of Post-Quantum Cryptography across Heterogeneous Computing Environments." *Cryptography*, vol. 9, no. 2, 2025, art. 32. 10.3390/cryptography9020032.

II. Al-Shetwi, Ali Q., et al. "Digital Twin Technology for Renewable Energy, Smart Grids, Energy Storage and V2G Integration." *IET Smart Grid*, vol. 8, no. 1, 2025. https://ietresearch.onlinelibrary.wiley.com/journal/25152947.

III. Bowers, James, and Mark Watson. "Constrained Device Performance Benchmarking with the Implementation of Post-Quantum Cryptography." *Cryptography*, vol. 8, no. 2, 2024, art. 21. 10.3390/cryptography8020021.

IV. Bundesamt für Sicherheit in der Informationstechnik. *Quantum-Safe Cryptography: Fundamentals, Current Developments and Recommendations*. BSI, 2024, https://www.bsi.bund.de.

V. Coppolino, Luigi, et al. "Exploiting Digital Twin Technology for Cybersecurity Monitoring in Smart Grids." *Proceedings of the 18th International Conference on Availability, Reliability and Security*, ACM, 2023. 10.1145/3600160.

VI. Cybersecurity and Infrastructure Security Agency, et al. *Quantum-Readiness Factsheet*. CISA, 2024. https://www.cisa.gov.

VII. Cybersecurity and Infrastructure Security Agency, et al. *Quantum-Readiness: Migration to Post-Quantum Cryptography*. Joint Advisory, CISA, 2023, https://www.cisa.gov.

VIII. Empl, Philip, et al. "Digital Twins in Security Operations: State of the Art and Future Perspectives." *ACM Computing Surveys*, vol. 58, no. 1, 2025. https://dl.acm.org/journal/csur.

IX. European Union Agency for Cybersecurity. *Post-Quantum Cryptography: Current State and Quantum Mitigation*. ENISA, 2022. https://www.enisa.europa.eu.

X. Hussain, S. M. Suhail, et al. "Performance Evaluation of IEC 61850 GOOSE Messaging for Time-Critical Substation Operations." *IEEE Access*, vol. 8, 2020. https://ieeexplore.ieee.org/xpl/RecentIssue.jsp?punumber=6287639.

XI. Institute of Electrical and Electronics Engineers. *IEEE 1815: Standard for Electric Power Systems Communications — Distributed Network Protocol (DNP3)*. IEEE, 2012. https://standards.ieee.org.

XII. Institute of Electrical and Electronics Engineers Power and Energy Society. *IEEE 13-Bus and 34-Bus Distribution Test Feeders*. Distribution Test Feeder Working Group. https://cmte.ieee.org/pes-testfeeders/.

XIII. International Electrotechnical Commission. *IEC 61850 Edition 2: Communication Networks and Systems for Power Utility Automation*. IEC, 2013. https://webstore.iec.ch.

XIV. International Electrotechnical Commission. *IEC 61850-5: Communication Requirements for Functions and Device Models*. IEC, 2013. https://webstore.iec.ch.

XV. International Electrotechnical Commission. *IEC 61850-8-1: Specific Communication Service Mapping — Mappings to MMS and to ISO/IEC 8802-3*. IEC, 2020, https://webstore.iec.ch.

XVI. International Electrotechnical Commission. *IEC 62351-3: Profiles Including TCP/IP*. IEC, 2018. https://webstore.iec.ch.

XVII. International Electrotechnical Commission. *IEC 62351-4: Profiles Including MMS and Derivatives*. IEC, 2018. https://webstore.iec.ch.

XVIII. International Electrotechnical Commission. *IEC 62351-6: Security for IEC 61850*. IEC, 2020. https://webstore.iec.ch.

XIX. International Electrotechnical Commission. *IEC 62351-9: Cyber Security Key Management for Power System Equipment*. IEC, 2023. https://webstore.iec.ch.

XX. Kannwischer, Matthias J., et al. *pqm4: Post-Quantum Crypto Library for the ARM Cortex-M4*. GitHub. https://github.com/mupq/pqm4.

XXI. Kniphoff da Cruz, Alexandre, et al. "IEC 61850 GOOSE: A Systematic Literature Review." *Automation*, vol. 7, no. 2, 2026, art. 62. 10.3390/automation7020062.

XXII. MZ Automation. *libiec61850: Open-Source IEC 61850 Library*. Version 1.5. https://libiec61850.com.

XXIII. National Institute of Standards and Technology. *Module-Lattice-Based Key-Encapsulation Mechanism Standard*. FIPS 203, NIST, 2024. 10.6028/NIST.FIPS.203.

XXIV. National Institute of Standards and Technology. *Module-Lattice-Based Digital Signature Standard*. FIPS 204, NIST, 2024. 10.6028/NIST.FIPS.204.

XXV. National Institute of Standards and Technology. *Stateless Hash-Based Digital Signature Standard*. FIPS 205, NIST, 2024. 10.6028/NIST.FIPS.205.

XXVI. National Institute of Standards and Technology. *Transition to Post-Quantum Cryptography Standards*. NIST IR 8547, Initial Public Draft, NIST, 2024. 10.6028/NIST.IR.8547.ipd.

XXVII. National Institute of Standards and Technology National Cybersecurity Center of Excellence. *Migration to Post-Quantum Cryptography*. NCCoE, 2024. https://www.nccoe.nist.gov.

XXVIII. National Security Agency. *Commercial National Security Algorithm Suite 2.0*. NSA, 2022. https://www.nsa.gov.

XXIX. Ofenloch, Annika, et al. "MOSAIK 3.0: Combining Time-Stepped and Discrete Event Simulation." *2022 Open Source Modelling and Simulation of Energy Systems (OSMSES)*, IEEE, 2022. 10.1109/OSMSES54027.2022.9769116.

XXX. Open Quantum Safe Project. *liboqs and oqs-provider for OpenSSL 3*. Open Quantum Safe, https://openquantumsafe.org.

XXXI. Qureshi, Ayesha, et al. "A Survey on Security-Enhancing Digital Twins." *Computer Communications*, vol. 238, 2025. https://www.sciencedirect.com/journal/computer-communications.

XXXII. Reda, Haftu Tasew, et al. "Vulnerability and Impact Analysis of the IEC 61850 GOOSE Protocol in the Smart Grid." *Sensors*, vol. 21, no. 4, 2021, art. 1554. 10.3390/s21041554.

XXXIII. Sánchez, Guillermo, et al. "Masquerading the IEC 61850 GOOSE Protocol: Cyber-Physical Experiments and Detection." *Proceedings of the 16th ACM International Conference on Future Energy Systems*, ACM, 2025. https://dl.acm.org/conference/e-energy.

XXXIV. Sen, Ömer, et al. "Digital Twin for Evaluating Detective Countermeasures in Smart Grid Cybersecurity." *IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids*, IEEE, 2023. https://ieeexplore.ieee.org.

XXXV. Steinbrink, Cornelius, et al. "CPES Testing with Mosaik: Co-Simulation Planning, Execution and Analysis." *Applied Sciences*, vol. 9, no. 5, 2019, art. 923. 10.3390/app9050923.

XXXVI. Wang, Qian, and Zhen Li. "Upgrading Operational Technology Systems to Post-Quantum Cryptography: Strategies and Resource Models." *International Journal of Critical Infrastructure Protection*, vol. 42, 2023, art. 100612. 10.1016/j.ijcip.2023.100612.