JOURNAL OF POST-QUANTUM CRYPTOGRAPHY AND SYSTEMS

RESEARCH ARTICLE

Volume 01, Issue 02, 2026

An Adaptive Quantum-Resistant Intrusion Detection and Secure Communication Framework for Next-Generation Cyber-Physical Systems

Mahfuz Saeed Zada

Department of Computer Science and Engineering, Caspian Institute of Industrial Engineering, Iran

ABSTRACT

Next-generation cyber-physical systems need to be protected by security mechanisms that are able to adapt to new attacks and to withstand the new threats that are possible in the future with quantum technologies. This study designed an adaptive quantum-resistant intrusion-detection and secure-communication framework that combines the elements of behaviour monitoring, concept-drift adaptation, post-quantum key establishment, digital signature, policy control, and automatic response. CPS dataset with 1,500 observations and 38 variables was generated, along with 480 post-quantum cryptographic benchmark trials, for a simulation-based design. The CPS data covered normal traffic, 6 different categories of attacks, 94 zero-day attacks, and 218 concept-drift observations. The baseline classifiers were logistic regression, decision tree, random forest, SVM, and gradient boosting, which were evaluated, and the proposed adaptive IDS was compared with a static one. The adaptive model achieved 95.67% accuracy, 98.99% precision, 90.88% recall, a 94.76% F1-score, 99.35% ROC-AUC, and a 0.70% false-positive rate. The recall accuracy on zero-day improved from 28.72% to 77.66% and concept-drift adaptation accuracy improved from 55.96% to 81.65%. Post-quantum profiles offered more long-term security, but necessitated more handshake time, memory, energy, and communications overhead. The results indicated that the most efficient and quantum-resistant combination was the use of ML-KEM-768 and ML-DSA-65. The results show the potential of adaptive intrusion detection combined with post-quantum communication for next-generation CPS, but testing on physical testbeds and real operational datasets is required.

Keywords: Adaptive intrusion detection; cyber-physical systems; post-quantum cryptography; zero-day attacks; secure communication

1. Introduction

Next-generation cyber-physical systems combine computation, sensing, communication, control and physical processes to enable intelligent and autonomous operation. Environmental and operational data are constantly gathered by sensors, communicated via communication networks, computed through computational platforms, and acted upon physically by controllers.Sensors gather environmental and operational data, communication networks transmit data, computational platforms analyse conditions in the system, and controllers take physical action. These interrelated capabilities are used more in the fields of smart grid, autonomous transportation, industrial automation, connected healthcare and smart city infrastructure. But, due to the level of digital involvement, cyber incidents may have direct operational, economic and safety impacts. Thus, monitoring using physics has come into prominence to detect malicious activities that change normal system behaviour or physical-process variables (Giraldo et al., 2019). However, cooperation between security controls is needed for industrial cyber-physical systems, as a compromise in the network, control or device layer can impact the entire operational environment (Ding et al., 2018).

Malware, DoS attacks, false data injection, spoofing, replay attacks, command manipulation, insider operations and advanced persistent threats are among the various attack scenarios that pose threats to cyber-physical systems. An attacker can tamper with sensor readings, disrupt communications, masquerade as an authorized sensor or send unauthorized actuator commands. The datasets collected from modern IoT and industrial IoT scenarios show that these environments produce heterogeneous telemetry and network traffic, which need special data-driven detection methods (Alsaedi et al., 2020). Multi-layered detection has become a significant concern given the possibility of attacks being first visible in network traffic, and then later in control processes or physical operations. One way to enhance visibility across the interconnected industrial-control components is to implement a three-tiered intrusion detection structure (Anthi et al., 2021). Cognitive Computing has also been shown to have potential for detecting complex attack patterns in Industrial Cyber Physical Environments (ICPE) (Althobaiti et al., 2021).

RSA and elliptic-curve cryptography are often used for authentication, key exchange and digital signatures in existing CPS communication mechanisms. If quantum computer technology is advanced enough, these mechanisms could be compromised and data, devices and control instructions could be intercepted. Thus, post-quantum cryptography has become a critical technology for the secure protection of connected devices and IoT infrastructures in the long run (Fernández-Caramés, 2019). Through recent NIST evaluations, quantum-resistant algorithms have been evaluated progressively over the following criteria: security, implementation efficiency, and practical deployability (Alagic et al., 2020). The second step of standardisation further reduced the number of possible algorithms for secure post-quantum deployment (Alagic et al., 2022). There are quantum-resistant key establishment schemes based on lattices, like CRYSTALS-Kyber (Bos et al., 2018) and quantum-resistant digital signatures, like CRYSTALS-Dilithium (Ducas et al., 2018).

Traditional intrusion detection systems are generally signature-based, have pre-defined thresholds and static classification rules. These can identify previously seen attacks but have limited ability to deal with zero-day attacks, changing traffic behaviour and concept drift. The use of machine-learning and deep learning systems for feature learning is also better, but this approach requires representative datasets and selecting the right machine learning models (Ahmad et al., 2021). Although anomaly-based deep-learning models detect unknown attacks, they suffer from high false-alarm rates, limited interpretability and a high computational cost (Aldweesh et al., 2020). Comparative studies also show large differences between the performance of

models for different datasets and attack classes (Ferrag et al., 2020). While generative adversarial learning can enhance CPS intrusion detection, in distributed fog environments, there are extra training and resource demands to be addressed (de Araujo-Filho et al., 2020). Specialized systems like DEIDS show better performance in detecting attacks in industrial systems, but are only effective under certain operational conditions and feature distributions (Gu et al., 2022).

Protecting CPS needs more than just independent intrusion detection or encryption. Adaptive attack detection should be used in conjunction with quantum-resistant authentication, secure key establishment and encrypted data transmission. This integration can make sure that suspicious activities are detected in the communication process and that information from the sensors, commands and identities of the devices is kept secure.

The purpose of the study is to design and test an intelligent intrusion detection and secure communication-based adaptive quantum-resistant framework for next-generation CPS. The goals are to create a layered security design, identify known and zero-day attacks, apply the concept of post-quantum key establishment and authentication, analyse conventional and quantum attacks, compare the design with existing designs, analyze the computational efficiency, scalability, communication overhead and applicability in real time.

2. Materials and Methods

2.1 Research Design and System Architecture

An adaptive quantum-resistant security framework for next-generation cyber-physical systems was simulated and evaluated using experimental design. The architecture included smart grids, industrial automation, autonomous transportation, healthcare and smart city environments. It consisted of physical-process, sensing, communication, edge, control and application layers. Operational data was collected by sensors and devices, information was communicated via protocols, local processing was provided by edge gateways and control units acted on physical actions.

2.2 Proposed Adaptive Quantum-Resistant Framework

The framework included 5 components an intrusion-detection engine, an adaptive-learning unit, a post-quantum cryptographic module, a policy controller and an automated response mechanism. The adaptive unit calculated the model under concept drift and the detection engine computed the attack-probability scores. The cryptographic module ensured authentication, key exchange and data transfer. The response mechanism started monitoring, rate limiting, source blocking, device isolation or key revocation or model updating depending on the estimated security risk.

2.3 Threat Model and Attack Scenarios

The security risks considered were the ability to intercept network traffic, spoof legitimate nodes and alter sensor readings, replay messages, and to introduce commands or malware. There were 853 normal observations and 647 attack observations. Denial-of-service, false-data injection, malware, man-in-the-middle, spoofing and replay attacks were simulated attacks. It also had 94 zero-day attack cases and 218 concept drift observations to test adaptability to changing conditions.

2.4 Dataset Collection and Preprocessing

Two datasets were used in the study. There were 1,500 observations and 38 variables in the CPS data set, with each representing 5 minutes. The variables measured were network traffic, authentication failures, sensor deviation, control command activity, processor usage, memory usage, energy consumption, latency, throughput and packet delivery ratio. The cryptographic data set had 480 trials and 19 variables. There were no missing values in the datasets and there were no duplicate records. Categorical variables were one-hot encoded and continuous variables were scaled by using the training set parameters. The predictor set contained no identifiers, timestamps, prediction scores or mitigation actions; these were all excluded because they might be "leaked" into the outcome set. Non-zero-day records were split into a training subset of 70%, validation subset of 15% and testing subset of 15%. Zero-day records were left to be evaluated separately.

2.5 Adaptive Intrusion Detection Model

Logistic regression, decision tree, random forest and gradient boosting were selected as the baseline classifiers. The proposed adaptive IDS adopted network, authentication, process and behavioural characteristics and implemented continuous anomaly score calculation. Observations were classified as normal or malicious by using a threshold of 0.50. Chronological changes in feature distributions were used to assess concept drift. The adaptive model adjusted the decision boundary and the static IDS did not. The adaptive model adjusted the decision boundary when drift was detected while the static IDS did not. The zero-day performance was evaluated on attacks that were not trained.

2.6 Quantum-Resistant Secure Communication Module

Three modes of security were compared: classical, hybrid and full post-quantum. RSACerts 2048 and ECDSA P-256 were used in classical configuration. In the hybrid mode, RSA and ML-KEM were used together and ECDSA and ML-DSA were used together. Both fully post-quantum profiles (1024-bit, 768-bit, 512-bit) used ML-KEM-1024/ ML-DSA-1024, ML-KEM-768/ ML-DSA-768, or ML-KEM-512/ ML-DSA-512.

Another implementation of ML-KEM and SLH-DSA was tested. The key generation time, encryption/decryption time, signature generation time, signature verification time, handshake time, ciphertext size, signature size, memory usage and energy usage were measured for the benchmark dataset. Once key establishment was successful, authenticated symmetric encryption was assumed.

2.7 Experimental Setup and Implementation

The experiment was run in a software-based simulated environment not in real CPS hardware. Data processing, modelling and statistical analysis were done using Python-based tools. NumPy and pandas were used for processing, scikit-learn for machine-learning evaluation and SciPy and statsmodels for inferential testing. To increase reproducibility, a fixed random seed of 24072026 was used. Three levels of system load were tested: low, moderate, and high.

2.8 Performance Evaluation and Statistical Analysis

The accuracy, precision, recall, specificity, F1 score, false positive rate, false negative rate, ROC AUC and detection latency were used to assess intrusion-detection performance. Zero-day and concept-drift performances were analyzed separately. The adaptive and static IDS predictions were compared and analyzed by employing McNemar's test. Handshake time, memory usage, energy consumption, size of ciphertext, size of signature, network latency, throughput and packet delivery ratio were used to measure the cryptographic and system performance.

The Wilcoxon t-test, Mann–Whitney U test, one-way analysis of variance (ANOVA), chi-square analysis and Pearson correlation tests were used as appropriate. The criterion of statistical significance was set at $p < 0.05$, and Cohen's d, Cramer's V, and eta squared were used for reporting effect sizes.

3. Results

3.1 Dataset and Feature Characteristics

There were 1,500 observations and 38 variables in the CPS dataset, and 480 benchmark trials and 19 variables in the dataset for post-quantum cryptographic (PQC) algorithms. Of the CPS records, 853 were in normal operation and 647 were in a cyberattack. Denial of service was the most prevalent attack type, false-data injection was next, followed by malware, man-in-the-middle and spoofing/replay attacks. The data also contained 94 zero-day attacks along with 218 observations over the concept drift scenario.

There were no missing data or duplicate records in either dataset. The one-hot encoding method was used for categorical variables and the normalization method was used for continuous variables with parameters trained on the training data. No record identifiers, no time stamps, no IDS scores and outcomes derived from the IDS, no mitigation actions were used as inputs to the models. The most important predictive features determined by the random forest model are packet-delivery ratio, network latency, sensor deviation, CPU utilization, signal entropy, energy consumption, failed authentication attempts and packet rate. Table 1 shows the distribution of normal and attack observations, individual attack categories, zero-day results, concept drift results and data-quality results.

Table 1. Dataset and attack-class characteristics

Dataset characteristic	Frequency	Percentage
Total CPS observations	1,500	100.0%
Normal observations	853	56.9%
Attack observations	647	43.1%
Denial-of-service	147	22.7% of attacks
False-data injection	133	20.6% of attacks
Malware	108	16.7% of attacks
Man-in-the-middle	96	14.8% of attacks
Spoofing	88	13.6% of attacks
Replay	75	11.6% of attacks
Zero-day attacks	94	14.5% of attacks

Concept-drift observations	218	14.5% of CPS records
Cryptographic benchmark trials	480	100.0%
Missing values	0	0.0%
Duplicate records	0	0.0%

3.2 Performance of Baseline Intrusion Detection Models

The non-zero day (NZD) records were sorted by date and split into 984 NZD training, 210 NZD validation and 212 NZD testing observations. After the model selection, the train and validation subsets were merged for final model estimation.

On chronological test set, the accuracy, F1-score and ROC-AUC of logistic regression, random forest and support vector machine were 100%. Gradient boosting produced 99.53% accuracy, 98.97% recall, 99.48% F1-score and 99.99% ROC-AUC. The decision tree achieved 99.06% accuracy, 97.94% recall, 98.96% F1-score and 99.47% ROC-AUC. These scores were very high and demonstrated a high separation between normal and malicious observations. The comparative performance of the five baseline intrusion-detection models is shown in Figure 1.

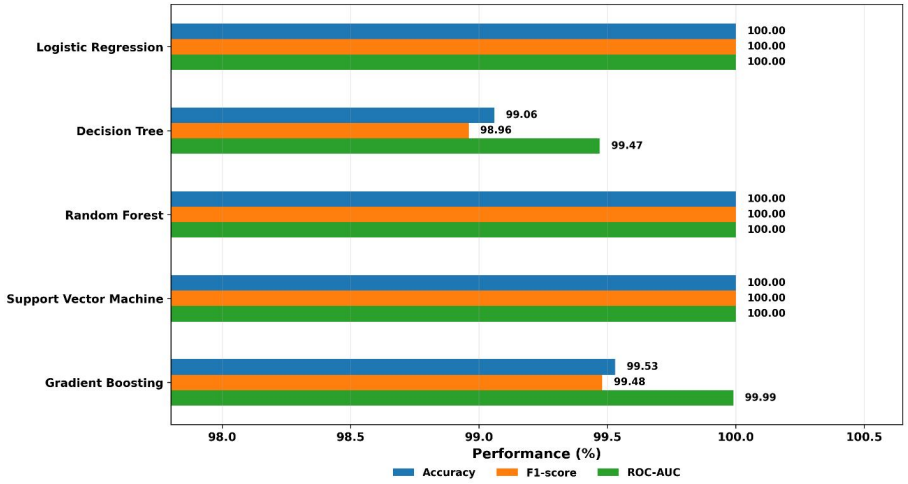

Figure 1. Performance of baseline intrusion-detection models

3.3 Performance of the Proposed Adaptive Intrusion Detection Model

The proposed adaptive IDS performed significantly better compared to the static IDS given, and this has been verified over all 1,500 observations. The adaptive model yielded accuracy of 95.67%, precision of 98.99%, recall of 90.88% and specificity of 99.30%. Its F1-score was 94.76%, while the ROC-AUC reached 99.35%. The false positive rate was constrained to 0.70%, whereas that of the static IDS was 4.57%.

The average detection latency of attack observations was 20.96 ± 10.69 ms and the median was 18.20 ms. The adaptive IDS correctly identified 220 observations that were incorrectly identified by the static IDS, while the static IDS only correctly identified 38 observations that were incorrectly identified by the adaptive IDS, as evidenced by a paired McNemar comparison. This difference was significant at $p < 0.001$ level. Detective performance in terms of accuracy, precision, recall, specificity, F1-score, false-positive rate, ROC-AUC and detection latency of the static and adaptive intrusion-detection systems is compared in Table 2.

Table 2. Performance of the static and adaptive intrusion-detection systems

Performance metric	Static IDS	Adaptive IDS
Accuracy	83.53%	95.67%
Precision	91.84%	98.99%
Recall	67.85%	90.88%
Specificity	95.43%	99.30%
F1-score	78.04%	94.76%
False-positive rate	4.57%	0.70%
ROC-AUC	92.49%	99.35%
Mean detection latency	Not available	20.96 ms
Median detection latency	Not available	18.20 ms

3.4 Detection of Known and Zero-Day Attacks

The adaptive IDS had a high recall in all six categories of attacks. The highest adaptive recall was 94.32 due to spoofing attacks, 91.67 due to malware, 91.16 due to denial of service, 90.67 due to replay, 90.63 due to man-in-the-middle and 87.97 due to false-data injection.

In the 553 known attacks, adaptive recall was found to be 93.13% as opposed to 74.50% with the static IDS. The difference in performance was even greater with the 94 zero-day attacks. The static IDS only detected 28.72% of these attacks which were not seen before compared to the adaptive IDS which detected 77.66%. The adaptive recall was thus about 49 percentage points greater in the zero-day conditions. The comparison of Figure 2 between static and adaptive recall of known, zero-day, concept-drift and complete attack subsets.

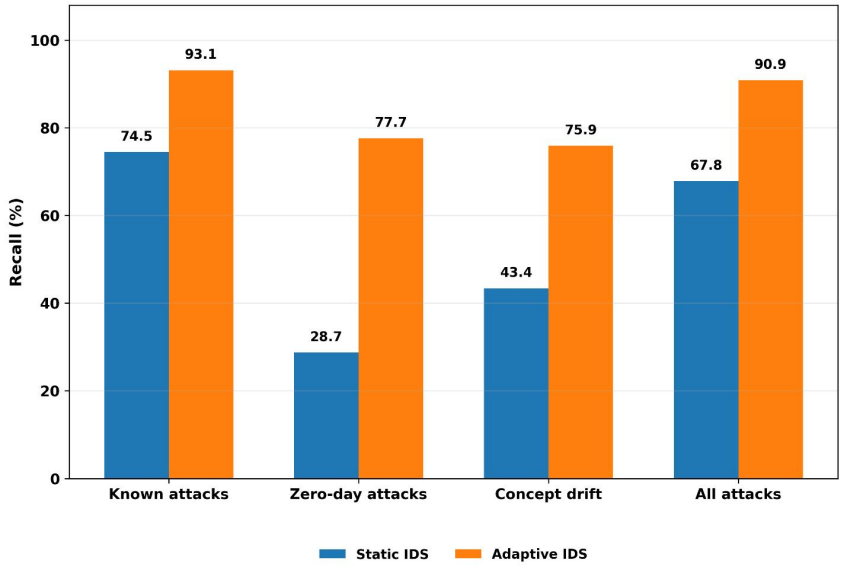

Figure 2. Static and adaptive IDS recall under known, zero-day and concept-drift conditions

3.5 Adaptability under Concept Drift

There were 218 concept-drift observations (with 166 attack observations). The model updating failed to produce high accuracy and recall figures as the static IDS only produced 55.96% and 43.37% accuracy and recall respectively. Following adaptive

updating, the accuracy and recall rose to 81.65 and 75.90, respectively.

The adaptive configuration was perfectly precise in the drift subset, which means that all observations that were labeled as attacks were truly malicious and no normal drift observations were labeled as such. The F1-score increased from 60.00% to 86.30%, while ROC-AUC increased from 85.61% to 98.59%. The test of McNemar proved the change after the adaptation to be statistically significant, $p < 0.001$. The results of adaptive model updating in concept-drift conditions are presented in Table 3 to demonstrate how intrusion-detection performance improves.

Table 3. IDS performance under concept-drift conditions

Performance metric	Static IDS without update	Adaptive IDS after update
Accuracy	55.96%	81.65%
Precision	97.30%	100.00%
Recall	43.37%	75.90%
Specificity	96.15%	100.00%
F1-score	60.00%	86.30%
False-positive rate	3.85%	0.00%
ROC-AUC	85.61%	98.59%

3.6 Post-Quantum Cryptographic Performance

There was a significant difference in cryptographic performance among the six security profiles, with all the studied operation times, ciphertext sizes, signature sizes, memory requirements and energy measurements yielding $p < 0.001$.

The highest mean secure-handshake time of 6.97 ms was recorded with the classical RSA-2048 and ECDSA profile. It had a mean key-generation, encapsulation, decapsulation, signature-generation and signature-verification time of 0.379, 0.240, 0.229, 0.596 and 0.747 ms, respectively.

The ML-KEM-512 and ML-DSA-44 profile took 8.89 ms to make a handshake, whereas ML-KEM-768 and ML-DSA-65 took 10.02 ms. The more secure ML-KEM-1024 and ML-DSA-87 setting took 12.43 ms. The hybrid RSA, ML-KEM-768 and ML-DSA-65 profile had a mean handshake time of 11.07 ms.

The profile created by SLH-DSA produced the largest overhead with the major reason being that its average signature-generation time was 18.19 ms. It had a total time of 28.80 ms in its handshake, as compared to 1.30 ms with signature generation and 10.02 ms with the ML-KEM-768 and ML-DSA-65 profiles.

The average ciphertext length grew to 255 bytes with the classical profile, and to 767 to 1,576 bytes with fully post-quantum profiles. The signature size doubled to 2,417-4,622 bytes with ECDSA to 7,870 bytes with ML-DSA and SLH-DSA respectively. The benchmark data set did not include public- and private-key sizes as individual variables and thus were not reported. Figure 3 shows the average secure-handshake time, and its variability between cryptographic profiles.

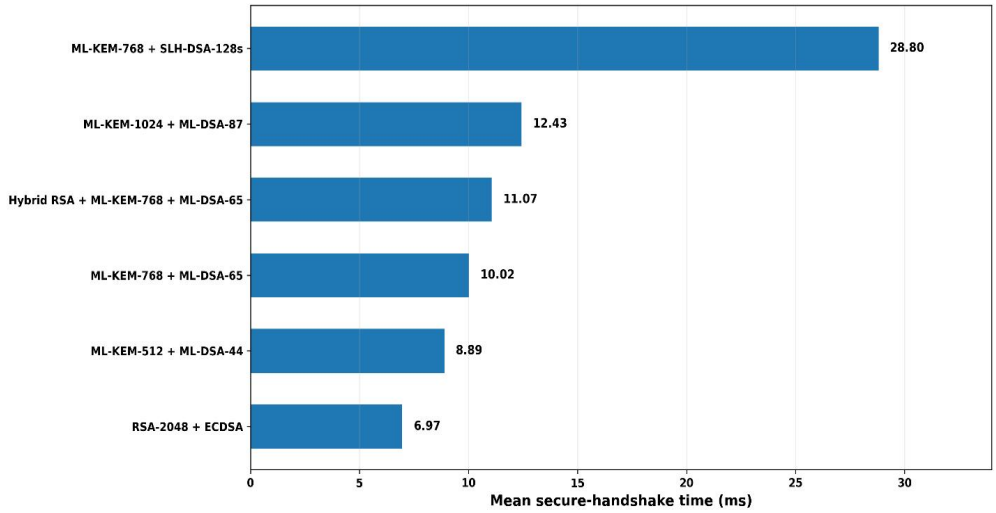

Figure 3. Secure-handshake performance across classical, hybrid and post-quantum cryptographic profiles

3.7 Communication and System-Level Performance

The cyberattacks resulted in significant degradation of the communication and resource performance in CPS. The mean network latency rose to 54.27 20.16 ms in the attacks compared to 23.48 7.08 ms under normal operation. Mean throughput dropped to 38.21 Mbps and 0.834 Mbps, whereas the packet-delivery ratio dropped to 0.834. The differences were statistically significant, $p < 0.001$.

The mean CPU utilization rose from 38.54 to 57.29, whereas memory utilization rose from 43.71 to 53.61. The consumption of energy rose to 3.91 J. The model response sizes were large in the latency, packet-delivery ratio, CPU utilization and energy consumption.

The lowest processing and communication overhead was the classical security mode. Hybrid and all-post-quantum operation added extra latency, CPU utilization, memory consumption and energy demands. The mean throughput, however, was quite comparable between classical operation and hybrid operation. Ciphertext overhead was the greatest in hybrid mode as it used classical and post-quantum cryptographic elements. Table 4 presents a summary of the performance of classical, hybrid post-quantum, and fully post-quantum security modes in terms of communication and computation.

Table 4. Communication and computational performance by security mode

Security mode	Latency (ms)	Throughput (Mbps)	CPU use (%)	Memory use (%)	Energy (J)	Ciphertext overhead (bytes)
Classical	31.75	48.02	42.66	46.21	3.12	255.68
Hybrid-PQ	36.54	48.32	45.66	47.67	3.44	1,440.91
Post-quantum	39.17	45.82	48.86	48.93	3.66	1,089.21

3.8 Comparative, Ablation and Scalability Results

Ablation comparison showed that, with adaptive updating disabled and the use of the fixed IDS, overall accuracy decreased by

12.14 percentage points, and F1-score decreased by 16.72 percentage points. Zero-day and concept-drift attacks incurred the greatest loss, which validates the claim that the adaptive-learning component played the greatest role in conditions that are dynamic and have never been encountered.

The cryptographic profiles continued to work with low, moderate and high system loads, but handshake time escalated markedly with system load, $F(2, 477) = 4.66$, $p = 0.010$. Peak memory consumption, $F(2, 477) = 12.07$, $p < 0.001$, and energy consumption, $F(2, 477) = 8.48$, $p < 0.001$, also increased significantly.

From low to high load, handshake time increased by 26.1% for the classical profile, 33.4% for the hybrid profile, 26.1% for ML-KEM-512 with ML-DSA-44, 22.6% for ML-KEM-768 with ML-DSA-65 and 37.5% for ML-KEM-1024 with ML-DSA-87. The SLH-DSA profile changed between 26.31 and 33.34 ms. Figure 4 shows the variation in the time of a secure handshake with changing system load.

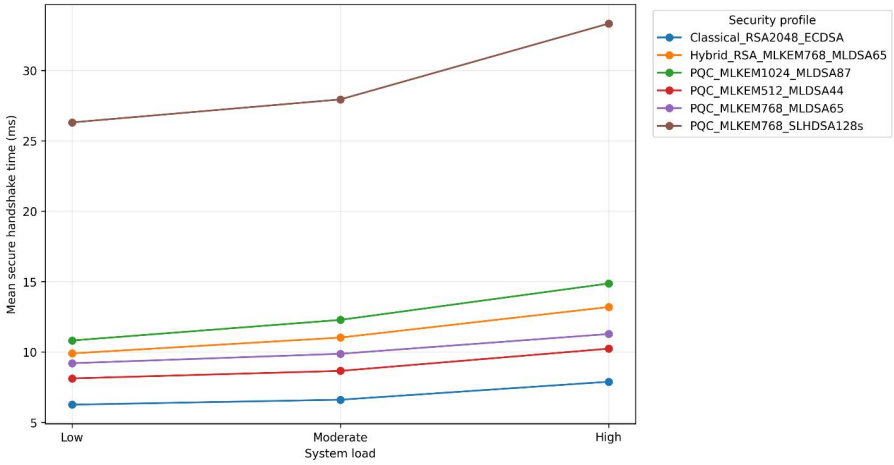

Figure 4. Scalability of classical, hybrid and post-quantum cryptographic profiles under increasing system load

ML-KEM-768 and ML-DSA-65 offered the best balance between fully post-quantum time and space, with a balance between the maximum quantum-resistance score of 95.29 and a mean time in the handshake of 10.02 ms. ML-KEM-1024 and ML-DSA-87 offered the highest quantum-resistance score. The dataset lacked a distinct variable of the number of devices hence scalability was measured based on system-load conditions as opposed to directly adding the number of connected CPS devices.

4. Discussion

The adaptive intrusion detection system was shown to be 95.67% accurate, precision 98.99, recall 90.88, and F1-score 94.76 which implies that the proposed system can distinguish between normal and malicious CPS activities. The network-flow, authentication, physical-process and device-resource variables were used together and this combination led to its low false-positive rate of 0.70%. These multidimensional characteristics helped the model to identify the real operational variations and those that were malicious. Past research also revealed that deep neural and machine-learning models are more effective in classifying attacks in the presence of representative behavioural features (Naseer et al., 2018; Zolanvari et al., 2019). Nevertheless, the very high baseline performance was also indicative of very high separation between attack recordings and

normal recordings. The composition of datasets and the representativeness of features can have a significant impact on the reported IDS accuracy and generalizability (Hindy et al., 2020).

The adaptive framework proved to be very effective in responding to familiar and new attacks. The recalls were 93.13% known and 77.66% zero-day attacks versus 74.50 and 28.72 respectively of the static system. This enhancement implies that anomaly-based behavioural analysis proved more useful than pre-defined attack signatures in the case of previously unknown patterns. Similar potential has been demonstrated by autoencoder-based systems, which learn the normal behaviour of a network and identify abnormalities without necessarily being provided with full attack signatures (Meidan et al., 2018). The worth of continuous anomaly scoring to adapt network environments is also exhibited by online ensembles like Kitsune (Mirsky et al., 2018). However, zero-day detection was still lower than known-attack detection, which proves that unseen behaviours are still a significant issue with IDS (Khraisat et al., 2019).

Monitoring concept drifts and updating models were found to be more accurate in drift conditions with 55.96-81.65 and 43.37-75.90 recall, respectively. These results suggest that the detection boundary could not become outdated due to the changing traffic and process distributions because of incremental adaptation. False alarms were also minimized with a dynamic adjustment of scores without loss in attack sensitivity. Domain-specific autoencoders have also been identified to enhance CPS detection through feature representation adaptation to specific operation environments (Thakur et al., 2021). Federated and distributed learning can also be useful in updating adaptively across various CPS locations without centralizing sensitive operational information (Li et al., 2020).

A combination of ML-KEM, ML-DSA and SLH-DSA enhanced the long-term security of device authentication, key generation and digital signatures. These systems minimize reliance on RSA and elliptic-curve systems that can be exposed to sufficiently strong quantum computers (Mosca, 2018). The hybrid architecture assisted in the gradual migration with classical and post-quantum protection. Also recent studies have explored how to make standardized post-quantum algorithms stronger with the help of quantum-generated randomness, but in practice, security is affected by the quality of implementation and entropy management (Chen, 2026).

Enhanced cryptographic security raised the handshake time, ciphertext size, memory usage and energy. Classical profiles took an average of 6.97 ms per handshake, and the completely post-quantum profiles took about 8.89-12.43 ms. The SLH-DSA setup had the worst latency due to its larger signatures and slower signature process. These findings show that quantum resistance provides overhead that is measurable yet typically manageable. Choosing to do so should then be based on CPS latency, energy and security requirements. Some of the performance penalties could be minimized by improving the randomness generation and cryptographic implementation (Chen, 2026).

The results are consistent with the existing data, which suggests machine-learning and deep-learning approaches perform better than traditional static detection in the industrial setting. Convolutional-based methods have managed to identify attacks on industrial control processes (Kravchik & Shabtai, 2018) and deep autoencoders have aided generic and domain-specific CPS monitoring (Thakur et al., 2021). Nevertheless, the existing IDS studies in the industrial domain still have some challenges related to bias in the datasets, the interpretability of the model, the cost of deployment and changing attack behaviour (Umer et al., 2022).

Smart-grid substations, industrial controllers, connected vehicles, medical monitoring devices, traffic systems and critical infrastructure gateways could be supported by the framework. Its adaptive IDS can detect evolving attack behaviour and post-quantum communication safeguards commands, sensor data and device identities. Local detection can be fast at deployment at edge gateways without the need for constant cloud connectivity. The framework should be tested on physical testbeds, heterogeneous devices and external ICS datasets to validate the framework in future studies. Further studies are needed to explore adversarially edited inputs, explainable detection, federated learning to protect privacy, lightweight post-quantum implementations and deployment in the field.

5. Conclusion

The research came up with a novel and tested framework of adaptive quantum-resistant intrusion-detection and secure-communication in next-generation cyber-physical systems. The architecture proposed was network and process monitoring, adaptive learning, post-quantum key establishment, digital signatures, policy control, and automated response. The results of the experimental study of the CPS dataset revealed that the adaptive intrusion-detection system has significantly better results than the fixed one. It achieved 95.67% accuracy, 98.99% precision, 90.88% recall, a 94.76% F1-score, and a false-positive rate of only 0.70%. It performed well especially on known attacks, whereas the zero-day recall rose by 28.72% using the static model and 77.66% using the adaptive model. In concept-drift situations, adaptive updating enhanced accuracy to 81.65 as compared to 55.96% which validated the worth of continuous monitoring and model updating. The post-quantum communication profiles enhanced the resilience to future quantum-enabled attacks, but they added to the handshake latency, ciphertext size, memory usage and energy usage. ML-KEM-768 (ML-DSA-65) turned out to be the most reasonable solution as a balance between quantum resistance and computational efficiency. Overall, it was shown that adaptive intrusion detection and post-quantum communication can be co-implemented without causing prohibitive overhead on the system. Nevertheless, the results were obtained based on artificial data and model measurements. The framework needs to be tested on physical testbeds, heterogeneous CPS devices, external industrial datasets, adversarial conditions, and long-term real-world deployments, which should be confirmed in the future.

References

1. Ahmad, Z., Shahid Khan, A., Wai Shiang, C., Abdullah, J., & Ahmad, F. (2021). Network intrusion detection system: A systematic study of machine learning and deep learning approaches. Transactions on Emerging Telecommunications Technologies, 32(1), e4150. https://doi.org/10.1002/ett.4150
2. Alagic, G., Alagic, G., Apon, D., Cooper, D., Dang, Q., Dang, T., Kelsey, J., Lichtinger, J., Liu, Y.-K., & Miller, C. (2022). Status report on the third round of the NIST post-quantum cryptography standardization process. https://csrc.nist.rip/external/nvlpubs.nist.gov/nistpubs/ir/2022/NIST.IR.8413.pdf
3. Alagic, G., Alperin-Sheriff, J., Apon, D., Cooper, D., Dang, Q., Kelsey, J., Liu, Y.-K., Miller, C., Moody, D., & Peralta, R. (2020). Status report on the second round of the NIST post-quantum cryptography standardization process. US Department of Commerce, NIST, 2, 69.

4. Aldweesh, A., Derhab, A., & Emam, A. Z. (2020). Deep learning approaches for anomaly-based intrusion detection systems: A survey, taxonomy, and open issues. Knowledge-Based Systems, 189, 105124.
5. Alsaedi, A., Moustafa, N., Tari, Z., Mahmood, A., & Anwar, A. (2020). TON_IoT telemetry dataset: A new generation dataset of IoT and IIoT for data-driven intrusion detection systems. Ieee Access, 8, 165130–165150.
6. Althobaiti, M. M., Kumar, K. P. M., Gupta, D., Kumar, S., & Mansour, R. F. (2021). An intelligent cognitive computing based intrusion detection for industrial cyber-physical systems. Measurement, 186, 110145.
7. Anthi, E., Williams, L., Burnap, P., & Jones, K. (2021). A three-tiered intrusion detection system for industrial control systems. Journal of Cybersecurity, 7(1), tyab006.
8. Bos, J., Ducas, L., Kiltz, E., Lepoint, T., Lyubashevsky, V., Schanck, J. M., Schwabe, P., Seiler, G., & Stehlé, D. (2018). CRYSTALS-Kyber: A CCA-secure module-lattice-based KEM. 2018 IEEE European Symposium on Security and Privacy (EuroS&P), 353–367. https://ieeexplore.ieee.org/abstract/document/8406610/
9. Chen, A. C. H. (2026). NIST Post-Quantum Cryptography Standard Algorithms Based on Quantum Random Number Generators. Journal of Web Engineering. https://doi.org/10.13052/jwe1540-9589.2555
10. de Araujo-Filho, P. F., Kaddoum, G., Campelo, D. R., Santos, A. G., Macêdo, D., & Zanchettin, C. (2020). Intrusion detection for cyber–physical systems using generative adversarial networks in fog environment. IEEE Internet of Things Journal, 8(8), 6247–6256.
11. Ding, D., Han, Q.-L., Xiang, Y., Ge, X., & Zhang, X.-M. (2018). A survey on security control and attack detection for industrial cyber-physical systems. Neurocomputing, 275, 1674–1683.
12. Ducas, L., Kiltz, E., Lepoint, T., Lyubashevsky, V., Schwabe, P., Seiler, G., & Stehlé, D. (2018). Crystals-dilithium: A lattice-based digital signature scheme. IACR Transactions on Cryptographic Hardware and Embedded Systems, 238–268.
13. Fernández-Caramés, T. M. (2019). From pre-quantum to post-quantum IoT security: A survey on quantum-resistant cryptosystems for the Internet of Things. IEEE Internet of Things Journal, 7(7), 6457–6480.
14. Ferrag, M. A., Maglaras, L., Moschoyiannis, S., & Janicke, H. (2020). Deep learning for cyber security intrusion detection: Approaches, datasets, and comparative study. Journal of Information Security and Applications, 50, 102419.
15. Giraldo, J., Urbina, D., Cardenas, A., Valente, J., Faisal, M., Ruths, J., Tippenhauer, N. O., Sandberg, H., & Candell, R. (2019). A Survey of Physics-Based Attack Detection in Cyber-Physical Systems. ACM Computing Surveys, 51(4), 1–36. https://doi.org/10.1145/3203245
16. Gu, H., Lai, Y., Wang, Y., Liu, J., Sun, M., & Mao, B. (2022). DEIDS: A novel intrusion detection system for industrial control systems. Neural Computing and Applications, 34(12), 9793–9811. https://doi.org/10.1007/s00521-022-06965-4
17. Hindy, H., Brosset, D., Bayne, E., Seeam, A., Tachtatzis, C., Atkinson, R., & Bellekens, X. (2020). A Taxonomy of Network Threats and the Effect of Current Datasets on Intrusion Detection Systems. IEEE

Access, 8, 104650–104675. https://doi.org/10.1109/ACCESS.2020.3000179

18. Khraisat, A., Gondal, I., Vamplew, P., & Kamruzzaman, J. (2019). Survey of intrusion detection systems: Techniques, datasets and challenges. Cybersecurity, 2(1), 20. https://doi.org/10.1186/s42400-019-0038-7
19. Kravchik, M., & Shabtai, A. (2018). Detecting Cyber Attacks in Industrial Control Systems Using Convolutional Neural Networks. Proceedings of the 2018 Workshop on Cyber-Physical Systems Security and PrivaCy, 72–83. https://doi.org/10.1145/3264888.3264896
20. Li, B., Wu, Y., Song, J., Lu, R., Li, T., & Zhao, L. (2020). DeepFed: Federated deep learning for intrusion detection in industrial cyber–physical systems. IEEE Transactions on Industrial Informatics, 17(8), 5615–5624.
21. Meidan, Y., Bohadana, M., Mathov, Y., Mirsky, Y., Shabtai, A., Breitenbacher, D., & Elovici, Y. (2018). N-baiot—Network-based detection of iot botnet attacks using deep autoencoders. IEEE Pervasive Computing, 17(3), 12–22.
22. Mirsky, Y., Doitshman, T., Elovici, Y., & Shabtai, A. (2018). Kitsune: An Ensemble of Autoencoders for Online Network Intrusion Detection (arXiv:1802.09089). arXiv. https://doi.org/10.48550/arXiv.1802.09089
23. Mosca, M. (2018). Cybersecurity in an era with quantum computers: Will we be ready? IEEE Security & Privacy, 16(5), 38–41.
24. Naseer, S., Saleem, Y., Khalid, S., Bashir, M. K., Han, J., Iqbal, M. M., & Han, K. (2018). Enhanced network anomaly detection based on deep neural networks. IEEE Access, 6, 48231–48246.
25. Thakur, S., Chakraborty, A., De, R., Kumar, N., & Sarkar, R. (2021). Intrusion detection in cyber-physical systems using a generic and domain specific deep autoencoder model. Computers & Electrical Engineering, 91, 107044.
26. Umer, M. A., Junejo, K. N., Jilani, M. T., & Mathur, A. P. (2022). Machine learning for intrusion detection in industrial control systems: Applications, challenges, and recommendations. International Journal of Critical Infrastructure Protection, 38, 100516.
27. Zolanvari, M., Teixeira, M. A., Gupta, L., Khan, K. M., & Jain, R. (2019). Machine learning-based network vulnerability analysis of industrial Internet of Things. IEEE Internet of Things Journal, 6(4), 6822–6834.