

Article

Hybrid Quantum–Classical Anomaly Detection for 5G Roaming Signalling with QKD/QSDC Support

Themba Ngobeni *  and Boniface Kabaso * 

Department of Information and Communication Technology, Faculty of Informatics and Design, Cape Peninsula University of Technology, District Six Campus, Cape Town 8000, South Africa

* Correspondence: 216187133@mycput.ac.za (T.N.); kabasob@cput.ac.za (B.K.)

Abstract

Signalling traffic in 5G Interconnect and Roaming Networks (IRNs) over the IP eXchange (IPX) faces man-in-the-middle (MITM) interception on N32 and GTP-U, billing fraud, and “Harvest Now, Decrypt Later” (HN DL) adversaries. This paper develops and validates qSiP, a hybrid quantum–classical framework integrating Quantum Key Distribution (QKD, BB84 decoy-state protocol) and Quantum Secure Direct Communication (QSDC, DL04) at OSI Layers 1–2 with a classical-plus-quantum ML classifier for signalling anomaly and billing-fraud detection. Evaluation spans NS-3, Mininet, and MicroK8s on a dual-PLMN Open5GS testbed with PacketRusher NB-IoT traffic and STRIDE-L threat modelling. Cryptographically, qSiP holds BB84 QBER within decoy-state thresholds, sustains key rates matched to N32 timing, and under HN DL conditions bounds adversary exposure to one key-rotation interval; N32 captures confirm 3GPP message format and handshake semantics end-to-end. For detection, the Hybrid configuration reaches 98.5–99.5% accuracy across environments (vs. 89.4–96.0% classical), reduces undetected billing fraud over four roaming paths, and adds under 5 ms latency per signalling exchange. Statistical tests confirm hybrid > classical at $p < 0.05$ across eMBB, uRLLC, mMTC, and BCE. The work contributes an integrated framework with simulation-based empirical validation in a standards-aligned 5G SA testbed and a methodological commitment for future quantum-roaming research: quantum ML detects but does not encrypt, while QKD/QSDC encrypt does not classify; these are two complementary, non-interchangeable roles.

Keywords: 5G; interconnect and roaming; IPX; NB-IoT; signalling anomaly detection; quantum key distribution; QKD; QSDC; hybrid quantum–classical machine learning; STRIDE-L



Academic Editor: Jose I. Moreno Novella

Received: 8 May 2026

Revised: 11 June 2026

Accepted: 13 June 2026

Published: 28 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

The signalling traffic using 5G Interconnect and Roaming Networks (IRNs) over the IP eXchange (IPX) is exposed to a stratified set of anomalies. The combination of man-in-the-middle (MITM) interception of N32 and GTP-U flows, billing-fraud injection across the path of Billing and Charging Evolution (BCE) and the future-quantum so-called Harvest Now, Decrypt Later (HN DL) adversary together further compromise the confidentiality, integrity, and operational integrity of inter-operator signalling [1–3]. The narrow bandwidth and predictable rate of change in NB-IoT signalling is a particularly appealing target to adversaries who can capture traffic now and decrypt it later when quantum capability becomes available [4,5]. The predictability is appealing for two complementary reasons. First, the regular cadence and small, stereotyped payload formats of NB-IoT signalling make traffic identification and selective capture straightforward, so an adversary mounting

a passive HNDL collection campaign can isolate the high-value control-plane stream from background traffic with minimal classification effort [6,7]. Second, the structural regularity of the captured ciphertext, uniform packet boundaries, repeating message templates, and bounded payload entropy provides crib-like statistical structure that simplifies later cryptanalysis once a sufficiently capable quantum decryptor becomes available, in much the same way that highly templated message structure simplified classical cryptanalysis of routine military and diplomatic traffic.

The single, operational, defensive question that this paper will answer is as follows: how then can these anomalies be detected reliably, in real time, on a standards-aligned 5G SA roaming path when the adversary is assumed to be combining a classical interception capability with future cryptanalysis capabilities in quantum form? The solution sought in this case is to mix two distinct quantum-related abilities. At OSI Layers 1 and 2, the information-theoretic secrecy of an underlying signalling channel is provided by Quantum Key Distribution (QKD) and Quantum Secure Direct Communication (QSDC). The second is analytical: a hybrid classical-plus-quantum machine learning classifier consumes signalling traces and classifies them as legitimate, anomalous or fraudulent. Most importantly, the two capabilities have complementary functions. QKD/QSDC encrypt; they do not classify. Quantum-enhanced ML classifiers learn and detect; they are not encrypting. This framework utilises both, but does not confuse them, a distinction we draw out at the outset because in the broader literature, the two roles are sometimes conflated. For example, several recent surveys of “quantum-enhanced 5G/6G security” group QKD-based key establishment and QML-based intrusion detection under a single “quantum security” capability without articulating that the two address different layers of the threat model (the secrecy of the channel versus the labelling of behaviour on it) [8]. Similarly, several application-domain studies that combine QKD and QML in a single architecture present the QML stage as part of the same “quantum-secure” pipeline rather than as a complementary detection layer, conflating cryptographic guarantees with statistical classification performance [9]. Drawing the distinction explicitly is one of the methodological contributions of the present paper.

1.1. Objective and Hypothesis

This paper has a single objective, deliberately scoped to match the empirical evidence reported here:

Develop and provide proof-of-concept validation, in a simulated environment, of a hybrid quantum–classical anomaly detection framework for 5G IRN signalling, integrating QKD/QSDC at L1/L2 with hybrid ML classification to detect signalling anomalies and billing fraud.

The null hypothesis (H_0) states that integrating QKD/QSDC with hybrid quantum–classical ML produces no measurable improvement, against classical-only baselines, in either the security or the operational performance of 5G IRN signalling. The alternative hypothesis (H_1) states that the integration of quantum cryptographic primitives (QKD and QSDC) and quantum-enhanced ML classifiers measurably improves the security and operational performance of 5G IRN signalling compared to classical-only baselines.

1.2. Contribution

There are three contributions made in this paper. First, it defines a hybrid quantum–classical system of anomaly detection that explicitly separates the cryptographic task of quantum methods (QKD/QSDC at L1/L2) and the analytical task (quantum-enhanced ML classification) and how the two interact. Second, it reports empirical evidence on three fronts which have not previously been brought together in a single 5G SA standards-aligned evaluation:

cryptographic security at L1/L2, including a key-rotation analysis under simulated HNDL adversary conditions; preservation of 3GPP N32 message-format and handshake semantics; and hybrid ML detection performance against classical baselines across four roaming scenarios and four service domains. Third, it makes explicit a methodological clarification that quantum ML classifiers detect attacks while QKD/QSDC encrypt signalling and that the two are not interchangeable.

1.3. 5G Roaming Signalling: Background for Non-Telecom Readers

This subsection provides minimal 5G roaming background for readers from neighbouring fields (cryptography, machine learning, internet security). Specialist readers may skip directly to Section 2. A mobile network operator (MNO) runs what 3GPP calls a Public Land Mobile Network (PLMN). When a subscriber travels outside their operator's coverage, they roam onto a Visited PLMN (VPMN) while remaining contractually attached to their Home PLMN (HPMN); the two operators carry the resulting signalling and user-plane traffic to each other across an inter-operator backbone called the IP eXchange (IPX), under commercial and technical agreements standardised by the GSM Association (GSMA) [3,10,11]. Three actors are involved in every roaming session: the user equipment (UE) that consumes services, the HPMN that owns the subscription and authorises usage, and the VPMN that provides local radio coverage. Roaming over IPX inherits a long history of signalling-protocol vulnerabilities (SS7, GTP, Diameter, and most recently HTTP/2-based 5G signalling), and the operational consequences of inter-operator MITM, billing fraud, and subscriber-location tracking are well documented in the ENISA threat landscape for signalling security [12] and in measurement studies of global roaming infrastructure [13]. Figure 1 below shows the roaming reference architectures.




Figure 1. 5G roaming reference architectures (adapted from GSMA NG.113 [11]). (a) Home-routed (HR): user-plane traffic is tunnelled from the visited UPF back to the home UPF over N9 (GTP-U) and

exits to the data network (DN) at the home operator on N6. **(b)** Local breakout (LBO): user-plane traffic exits at the visited UPF on N6 and does not traverse the IPX. In both scenarios, inter-PLMN control-plane signalling crosses the IPX over the N32 interface between vSEPP and hSEPP—the surface protected by qSiP in this paper. Service-based interfaces (Nnssf, Nnrf, Npcf, Nudm, Nausf, Naf, Nsmf, Namf, Nsmf) connect the network functions within each PLMN; the dashed vertical line marks the inter-operator boundary.

In Figure 1, 5G Standalone (SA), the inter-PLMN signalling boundary is mediated by the Security Edge Protection Proxy (SEPP), a 3GPP-defined gateway at each operator's edge that protects the inter-operator interface as specified in TS 33.501 [1] and TS 29.573 [2]. SEPP-to-SEPP signalling exchanges run over an interface named N32, which has two logical sides: N32-c carries the control-plane handshake that the two SEPPs use to establish a secure context, and N32-f carries the forwarded operational signalling messages between them. N32-f messages are JSON-encoded and protected with the JOSE/JWE envelope (a signed, encrypted JSON wrapper), so a typical roaming message is structurally a JSON document inside an encrypted envelope inside a TLS-protected HTTP/2 session. User-plane subscriber data takes a separate path: in home-routed (HR) scenarios it is tunnelled back to the HPMN via the GPRS Tunnelling Protocol—User Plane (GTP-U) on the N9 interface, while in local-breakout (LBO) scenarios it exits at the visited operator's UPF on the N6 interface. The radio profile we use in this study, Narrowband IoT (NB-IoT), is a 5G low-bandwidth radio designed for sensors, metres, and other low-data-rate machine-to-machine endpoints; its narrow bandwidth and predictable cadence make it attractive both to legitimate IoT deployments and, as discussed in Section 1, to passive HNDL adversaries. Finally, 3GPP partitions the service space into four service domains that recur throughout this paper: enhanced Mobile Broadband (eMBB) for sustained-throughput consumer traffic, Ultra-Reliable Low-Latency Communications (uRLLC) for sub-10 ms control loops, massive Machine-Type Communications (mMTC) for dense IoT signalling, and the Billing and Charging Element (BCE) for the integrity-critical roaming-charging path. Together, these elements (PLMN, IPX, SEPP, N32, GTP-U, NB-IoT, and the four service domains) define the surface on which the framework in this paper operates.

1.4. Paper Structure

Section 2 reviews relevant work on 5G roaming security, QKD/QSDC, hybrid quantum–classical ML, and the methodological position of this paper. Section 3 describes the approach taken. Section 4 specifies the framework. Section 5 describes the experimental testbed and evaluation protocol. Section 6 reports cryptographic-security results first (Sections 6.1–6.3): L1/L2 metrics, HNDL key-rotation analysis, and N32 semantic preservation. Section 6.4 onward reports anomaly-detection results: hybrid ML performance against classical and quantum-only baselines across three environments. Section 7 reports statistical analysis. Section 8 discusses validity, limitations, and the cryptographic/analytical distinction. Section 9 concludes.

2. Related Work

2.1. Security Vulnerabilities in 5G Roaming and IPX

The 5G IRN design uses Security Edge Protection Proxies (SEPPs) at the perimeter of PLMN to mediate inter-operators signalling exchanges across the IPX; the control-plane handshakes are configured to run on N32-c and the forwarding roaming signalling messages on N32-f [14]. In home-routed settings: Different home-routed deployment settings involve tunnelling user-plane traffic over GTP-U on the N9 interface in case of a home-routed scenario and a home-routed exit at the visited UPF in case of a local-breakout deployment settings. Although transport layer security (TLS) safeguards the SEPP-to-SEPP

transmission, the forwarding route through the IPX and the user-plane tunnelling display known weaknesses to MITM interception, GTP-U injection, and billing-fraud manipulation and, prospectively, quantum cryptanalysis in the HNDL adversary model [4,15].

2.2. Quantum Key Distribution and Direct Communication

The canonical QKD protocol is still considered to be BB84, introduced by Bennett and Brassard in 1984 and later extended by the decoy-state variants of Hwang, Lo, and others; information-theoretic secrecy is practically achievable on noisy quantum channels by allowing legitimate parties to bound the wiretapper's information from observed bit-error statistics [16,17]. Introduced by Deng and Long in 2004, and extending the philosophy of QKD, DL04 transmits the secret message itself through entangled or single-photon channels, with overt integrity checks against wiretapping attacks [18,19]. The two protocols interface different levels of a roaming signalling stack: BB84 at the key-establishment layer and DL04 at the direct-message layer, which we use to implement per-frame integrity at OSI L2. Two points of clarification are warranted on the L2 use of DL04. First, the bulk of the 3GPP signalling payload itself continues to traverse the classical N32 path encapsulated in the standard JOSE/JWE envelope; full QSDC transmission of every signalling byte would be neither bandwidth-efficient at current QSDC rates nor compatible with legacy SEPP interoperability. Second, DL04 in qSiP is therefore used not as the primary message-transport channel but as a per-frame integrity-and-eavesdropping-detection primitive: a short DL04-protected tag is computed and verified per L2 frame, and any tampering or eavesdropping along the classical path is detected through the DL04 integrity check (with the BB84-established ephemeral key seeding the DL04 exchange). This narrow but explicit role is what allows the integrity-check overhead to remain fixed at 175–180 μs /frame regardless of attack conditions as per (Table 1).

Table 1. Cryptographic security metrics at OSI L1/L2 across the four roaming scenarios. QBER values exceeding the decoy-state threshold of 3.5% trigger an explicit key-rotation event under qSiP, bounding adversary exposure to a single rotation interval (see Section 6.2).

Metric	Nominal (HR)	Nominal (LBO)	MITM (HR)	MITM (LBO)
BB84 QBER (%)	2.1	1.8	6.4 $\rightarrow$ rotation-triggered	5.9 $\rightarrow$ rotation-triggered
Decoy-state threshold (%)	3.5	3.5	3.5	3.5
Key generation rate (kbps)	1.42	1.51	0.98	1.02
Key-establishment latency (ms)	3.8	3.6	4.7	4.4
DL04 integrity-check overhead (μs /frame)	180	175	180	175
Total qSiP overhead per signalling exchange (ms)	4.0	3.8	4.9	4.6

2.3. Hybrid Classical/Quantum Machine Learning for Network Security

Hybrid quantum–classical classifiers combine the low-latency assets of classical models (SVM, KNN, neural networks) with the higher-resolution decision bounds which quantum kernel methods (QSVM, QkNN, QNN) can provide on small, structured feature sets. Recent work has demonstrated that quantum kernels built using the ZZ feature map can outperform classical baselines on traffic-classification tasks in which the feature space is of moderately high-dimensionality, but the training set is limited [20,21]. The hybrid architecture in this case routes most of the unambiguous traffic across classical models to provide low latency and reserves quantum models to ambiguous or border-line cases in which the classification certainty margin is small.

2.4. Position of This Paper

This paper deliberately scopes its contribution to the empirical question of hybrid quantum–classical anomaly detection in 5G roaming signalling, with QKD/QSDC as the cryptographic substrate. The methodological framework on which the design and evaluation are based is an integrated DSRM–CkDT framework, summarised here only insofar as it affects the experimental design. The aim of this scoping is to ensure the current paper focuses on a single, testable hypothesis and that the reported results align with that hypothesis. With respect to existing hybrid quantum–classical ML studies for network security and anomaly detection, the novelty of this work is threefold. First, it is to our knowledge the first study to evaluate a hybrid quantum–classical anomaly detector specifically on the 5G IRN/IPX N32 signalling path, where prior hybrid quantum ML work has focused either on generic intrusion-detection benchmarks (KDD-Cup, CICIDS, UNSW-NB15) [8] or on radio-access-network side-channel features [9]; neither setting exercises the 3GPP TS 29.573 JOSE/JWE message structure or the SEPP-mediated inter-PLMN handshake that constitutes the actual signalling attack surface for billing fraud and HNDL collection. Second, this paper integrates the hybrid ML detector with a per-frame DL04 QSDC integrity check and a BB84 QKD key-rotation policy on the same signalling path, while previous studies typically evaluate hybrid ML and quantum cryptography as separate contributions in adjacent papers. Third, the paper makes the cryptographic-vs-analytical role distinction explicit as a methodological commitment, addressing the conflation observed in surveys [8] and combined-pipeline studies [9] noted in Section 1. With respect to recent quantum-safe roaming studies, the closest prior work is the TLS-1.3 cross-border roaming security and performance analysis of Lastre et al. [22], which evaluates classical post-quantum-ready handshakes on the 5G roaming path but does not introduce QKD/QSDC primitives or hybrid ML detection. More recently, GSMA has begun formalising its own industry-standards approach to post-quantum 5G roaming security in the PQ.05 (2025) framework [23], which addresses the migration of SEPP-mediated signalling to post-quantum–classical cryptography rather than to quantum-channel primitives. The present paper occupies a complementary point in the same design space: it evaluates QKD and QSDC quantum primitives empirically on the N32 signalling path, contributing information-theoretically secure key establishment that does not depend on the computational hardness assumptions underpinning PQC. The recent deployment-focused literature on practical QKD/QSDC integration into 5G/6G [24,25] has highlighted three deployment limitations: limited QKD key generation rates relative to 5G data rates, the absence of standardised QKD–5G interfaces beyond ITU-T Y.3800/Y.3801, and the operational cost of maintaining quantum channels alongside classical channels. The framework presented here addresses the first limitation through the explicit per-session key-rotation policy of Section 6.2 (which decouples session secrecy from raw key throughput) and identifies the latter two as scope items for future work.

Why N32 is a distinct and challenging target for hybrid quantum ML. Three structural features of the N32 path distinguish it from the generic benchmarks on which prior hybrid quantum ML work has been evaluated. First, the JOSE/JWE envelope means the detector cannot inspect payload content under encryption without breaking the SEPP security model; the available features are confined to JOSE header metadata, HTTP/2 framing, and transport timing—a far more constrained space than the rich packet-content features of KDD-Cup, CICIDS, or UNSW-NB15. The ZZ feature map in qSiM’s QSVM component was chosen precisely because quantum kernels perform well on such small, structured feature sets, where classical kernels tend to plateau. Second, the SEPP-mediated handshake (N32-c) interleaves authentication with the anomaly-bearing operational messages (N32-f) on the same HTTP/2 connection; a packet-wise classifier that ignores handshake state

confuses authentication events with attack onsets, so qSiM ingests the full session trace and encodes the handshake phase as a feature dimension. Third, billing fraud on the BCE path exploits the semantic content of charging records (unit counts, session identifiers, roaming-path descriptors) rather than raw traffic volume, giving it a statistical signature unlike the volume-based attacks that dominate generic label sets and requiring the dedicated BCE-domain training and evaluation that qSiM’s four-domain design provides.

It is also important to quantify why control-plane keying is tractable where user-plane keying is not. At the measured NB-IoT signalling load of approximately 1.8 kbps on N32, the BB84 key generation rate of 1.4–1.5 kbps is sufficient to maintain per-session key rotation. By contrast, the 5G user plane may carry traffic at multi-gigabit rates—six to seven orders of magnitude beyond current QKD key generation capacity—so the per-session key-rotation strategy that decouples session secrecy from raw key throughput applies to the control-plane signalling protected by qSiP and does not extend to bulk user-plane encryption.

3. Approach

This paper is a results paper, with one hypothesis. It combines Design Science Research Methodology (DSRM) and Concept–Knowledge Design Theory (CkDT) [26–28] to draw the methodological apparatus that it needs. Only the parts of the approach that directly relate to the repeatability of the empirical results are reported here.

3.1. Design Sub-Cycles

The iterative build–evaluate work was organised into four sub-cycles. In Sub-cycle 1, a baseline prototype was established which is based on classical ML-based traffic classification and a simple BB84 key exchange protocol. In Sub-cycle 2, the newly added quantum kernel ML classifiers (QSVM, QkNN, QNN) and the DL04 QSDC layer at L2 were introduced. Sub-cycle 3 adjusted the cryptographic parameters (decoy-state intensity ratios, key-rotation interval, integrity-check density) as a result of feedback on their evaluation and included an explicit HNDL key-rotation analysis reported in Section 6.2. Sub-cycle 4 integrated the artefact within the dual-PLMN testbed and conducted the cross-environment evaluation.

3.2. Five-Stage Validation Strategy

There were five phases of validation. The 5G NR L1/L2 traffic was generated in MATLAB/Simulink v24.2, and classical ML was benchmarked in phase 1 (2022 Q1–2023 Q3). In phase 2 (2023 Q4–2024 Q3), NS-3 packet-level simulations were conducted to assess quantum-hybrid classifiers. In phase 3 (Q4 2024), Mininet was validated with control-plane logic in the four domains of 3GPP services. The prototype was used in a pilot using Kubernetes/Docker in phase 4 (2025 Q1). In phase 5 (2025 Q2), metrics were brought together and added to a single analytics dashboard as shown in Figures 2 and 3.




Figure 2. qSiM and qSiP validation roadmap.




Figure 3. qSiM and qSiP progressive validation strategy.

The roadmap in Figure 2, traces the five validation phases and the specific tools used in each: phase 1 (2022 Q1–2023 Q3)—MATLAB R2024b/Simulink for 5G NR L1/L2 traffic generation and classical-ML benchmarking; phase 2 (2023 Q4–2024 Q3)—NS-3 v3.44 for packet-level simulation of quantum-hybrid classifiers; phase 3 (2024 Q4)—Mininet with ONOS/OVS for control-plane validation across the four 3GPP service domains; phase 4 (2025 Q1)—MicroK8s/Docker pilot deployment with Vagrant-driven Open5GS dual-PLMN cores; phase 5 (2025 Q2)—unified analytics dashboard consolidating metrics across the prior phases. Qiskit Aer and SQUANCH are used throughout phases 2–5 for quantum circuit simulation and QSDC channel emulation; Wireshark/tshark are used throughout for packet-level verification.

Each phase in Figure 3, builds on the artefacts of the prior phase and uses the toolchain summarised in the caption of Figure 2 (MATLAB/Simulink → NS-3 → Mininet + ONOS/OVS → MicroK8s/Docker → unified analytics). The progression deliberately moves from controlled simulation toward more realistic deployment environments, allowing each phase to surface a different class of issues: PHY/MAC behaviour in MATLAB/Simulink, packet-level dynamics in NS-3, control-plane logic across service domains in Mininet, and containerisation-induced execution variability in MicroK8s/Docker.

4. Framework qSiM Model and qSiP Protocol Architecture

4.1. High-Level Architecture

The qSiM model consists of nine modules including the ingestion of NB-IoT signalling information, classification and if necessary cryptographic security of the signalling traffic between two PLMNs over the IPX. There are two architectural commitments that are key. The cryptographic and analytical layers are physically separated: The cryptographic protocol method (BB84 + DL04) is physically implemented at the SEPP layer to protect the channel, and the traffic features are fed into the hybrid ML for detecting. Second, the output of the detection module(s) are all used in the response module which processes the output of the classifiers and modifies the cryptographic refresh policy as appropriate.

As illustrated in Figure 4: NFs (module (1)) AMF, SMF, AUSF, UDM, PCF and UPF deployed in both HPLMN and VPLMN are points of entry and exit of NB-IoT signalling traffic. The Signal Traffic Collection (2) module groups local-breakout and home-routed streams. The fundamental Next-Generation IRN-IPX qSiM engine (3) manages real-time security and classification using the Roaming Traffic Ingestion and Hybrid ML Classification modules, which use classical models (SVM, KNN, NN) with low-ambiguity traffic and elevate ambiguous or high-risk traffic to quantum-enhanced classifiers (QSVM, QkNN, QNN). The Adversary Traffic Input (4) and Adversary Threats/Attacks (5) modules explicitly capture the adversarial traffic, and the cryptographic layer is implemented by the Quantum Secure-Signalling Protocol (qSiP, module (6)). The Compare and Analyse Response module (7) decides on automatic response actions, the qSiM Errors and Resolvers module (8) resolves anomalies in run-time, and all flows that have been successfully processed become Secured/Flagged Traffic outputs (9).




Figure 4. High-level conceptual architecture: a nine-module pipeline (modules 1–9, described in the text) integrating cryptographic protection (qSiP, Module 6: BB84 QKD + DL04 QSDC) and analytical detection (qSiM, Module 3: hybrid classical + quantum ML across the eMBB, uRLLC, mMTC, and BCE service domains). Solid black arrows show the primary signalling/data flow, including the N9-UPF user-plane ingress and the N32-c/N32-f control-plane path; hollow arrows show internal flow within the qSiM engine from the service lanes to the Hybrid ML Classification block; double-headed arrows show bidirectional interaction with the qSiP layer, the adversary inputs (4, 5), and the Compare-and-Analyse (7) and Errors-and-Resolver (8) modules; the looping arrow labelled “NG: IRN-IPX” marks the iterative error-resolution feedback, and the “Fake NF” arrow marks adversary-injected traffic. Design principles—Secure, Dynamic, Integrated, and Adaptable—are noted at the base.

4.2. qSiP Cryptographic Layer (BB84 + DL04)

Before detailing the cryptographic construction, we state its scope precisely. The qSiP layer operates exclusively on the 5G inter-PLMN control plane—the N32 interface between the vSEPP and hSEPP—and provides information-theoretically secure key establishment and per-frame integrity checking on that signalling path only. The user-plane subscriber data carried over GTP-U on the N9 interface (home-routed) or exiting at the visited UPF on N6 (local-breakout) is not touched by qSiP and continues to rely on existing 3GPP TLS protection; no claim is made that quantum key material is applied to user-plane traffic.

The qSiP layer is embedded into the SEPP layer and functions on the signalling messages which pass through the IPX between the home and visited PLMNs, using N32 signals. The DL04 enhancement to BB84 decoy-state QKD provides ephemeral symmetric keys for each signalling session and provides additional protection for direct secure communication on L2 per-frame signalling. The unified quantum circuit in Figure 4 is employed for both BB84 key generation and DL04 secure-message exchange: the bit-error statistics of the legitimate parties are continuously monitored, and the presence of interception is detected. The circuit operates in four stages. First, EPR pair generation via Hadamard and CNOT gates on qubit pairs ($q_0 \rightarrow q_1$), ($q_2 \rightarrow q_3$), ($q_4 \rightarrow q_5$) creates three Bell pairs for BB84 key generation. Second, measurement of the “home” halves (q_0, q_2, q_4) and “check” halves (q_1, q_3, q_5) into classical registers yields sifted key bits and decoy-state error rates. Third, QSDC teleportation applies a Pauli-Y gate on data qubit q_3 to encode one bit of secret data, which is then teleported through the remaining entangled pair via CNOT and Hadamard gates. Fourth, data readout via final measurements of q_2 and q_5 projects the teleported bits into the classical data registers (data_bit0, data_bit1). Modular stages are separated by barriers (grey dashed lines), enabling alternative QKD/QSDC primitives to be substituted.

This single unified circuit simultaneously yields a secure shared key and a one-time-pad encrypted data stream.

DL04 integrity mechanism: quantum-mechanical basis. Because the per-frame DL04 check is central to the cryptographic substrate, we set out its mechanism explicitly. BB84 decoy-state QKD produces an ephemeral symmetric key K_e for each signalling session. K_e is information-theoretically secure in the precise sense that the mutual information between an eavesdropper's measurements and K_e is exponentially small in the privacy-amplification output length, provided the observed QBER does not exceed the decoy-state threshold; K_e is generated without reliance on any computational-hardness assumption.

The DL04 layer then uses K_e to seed a per-frame tag $T = f(K_e, \text{frame_index}, \text{message_content})$, where f is a universal hash applied to the frame content and a session-specific sequence number. Because K_e is information-theoretically secure and a one-time pad provides perfect secrecy (Shannon) when the key is as long as the message and used only once, the tag inherits an information-theoretic integrity guarantee: an adversary with unbounded computation who does not know K_e cannot forge a valid tag T' for a tampered frame M' with probability greater than $2 \times 10^{-\ell}$, where ℓ is the tag length in bits.

Eavesdropping detection operates through a separate channel. An adversary who intercepts or replays N32-f frames without compromising the BB84 quantum channel does not possess K_e and therefore cannot compute a valid tag for an injected or modified frame; the hSEPP detects the mismatch and raises an anomaly event in the response module (Module 7). A short DL04-protected acknowledgement, also seeded by K_e , is exchanged out-of-band, so active tampering with the acknowledgement itself is detectable through the DL04 eavesdropping signature (anomalous photon statistics on the quantum sub-channel). The information-theoretic property of the integrity check thus derives from the one-time-pad construction over the BB84 key, while DL04 contributes an independent second eavesdropping-detection channel rather than serving as the primary message-transport path.

4.3. Hybrid ML Detection Layer

In Module 3, the pipeline is a hybrid one that has both classical + quantum ML. Low-latency models such as SVM, KNN, and NN are used for most clear classes, while models that are enhanced by quantum circuits (QSVM with the ZZ feature map, QkNN, QNN) are used for ambiguous or borderline cases. The classifier models are trained with labelled NB-IoT traces (packet size, latency, source/destination, modulation features extracted from MATLAB/Simulink OFDM/QAM simulations) and used to classify a legitimate NB-IoT flow with a flow from an adversary in the four 3GPP service domains. Note: the hybrid ML layer is used to detect anomalies, but it is NOT used to encrypt the signalling. The encryption is provided by the qSiP cryptographic layer. The overall qSiM performance shown in Figure 4 is captured by a weighted utility function (Equation (1)):

$$U_{\text{qSiM}} = \omega_1(R_{\text{meas}}/R_{\text{max}}) + \omega_2 \cdot F_{\text{tele,meas}} + \omega_3 \cdot \text{Acc}_{\text{hybrid}} + \omega_4(1 - \text{PL}) - \omega_5 \cdot E_{\mu} \quad (1)$$

where R_{meas} is the measured secure key rate normalised by its theoretical maximum R_{max} , $F_{\text{tele,meas}}$ is the teleportation fidelity on the data register, $\text{Acc}_{\text{hybrid}}$ is the hybrid classifier accuracy, PL is the packet-loss rate, and E_{μ} is the measured QBER (decoy-state error). The weights $\omega_1, \dots, \omega_5$ are normalised to sum to unity; the values used in the present evaluation are $\omega = (0.25, 0.25, 0.30, 0.10, 0.10)$.

The hybrid classification probability combines classical and quantum outputs via a mixing weight $\lambda \in [0, 1]$ (Equation (2)):

$$P_{\text{combined}}(y | x) = (1 - \lambda) \cdot P_{\text{classical}}(y | x) + \lambda \cdot P_{\text{quantum}}(y | x) \quad (2)$$

with λ tuned on the held-out validation set (the value used in this paper is $\lambda = [0.5]$). The qSiP one-time-pad encryption operates as (Equation (3)):

$$C = M \oplus K \quad (3)$$

with decryption identical to encryption.

Quantum Kernel Classifier Specifications

The quantum components of the analytical layer are chosen to be sufficiently small to be able to be implemented in statevector-simulator hardware; they are structurally and functionally independent of the qSiP cryptographic circuit in Figure 5. Note that the qSiP cryptographic circuit (Figure 5) has 6 qubits to support three Bell pairs (check, data, and home registers); the analytical-layer circuits here are independent qubit budgets of the size of the dimensionality of the features of interest in the classification task, with 4 qubits being sufficient to amplitude-encode the four major NB-IoT signalling features listed (packet size, inter-arrival time, source/destination pair, modulation class). QSVM applies a ZZ feature map of 4 qubits and two repetitions; the classically defined feature inputs (size of packet, can be interarrived by second packet, source/destination pair) are amplitude-encoded into the feature map. QkNN involves a swap-test circuit that has 5 qubits, and in which the test ancilla controls the measurement of the overlap between the query state and the reference state. QNN is a 4-qubit hardware-efficient ansatz based on a single layer of parameterised RY rotation, and a chain of CNOTs. The circuits have no more than 10 qubits, which is in line with the simulator envelope in this research.




Figure 5. qSiP unified quantum circuit (QSDC-DL04 with BB84 key seeding), drawn in standard Qiskit circuit notation. The horizontal lines q_0 – q_5 are qubits and the lower lines (check, home, data_bit0, data_bit1) are classical registers, the slashed number giving each register's bit-width. Red boxes marked H are Hadamard gates; a blue control dot joined to a blue $\oplus$ is a CNOT gate; the maroon box marked Y is a Pauli-Y gate; grey boxes with a meter symbol are measurements, whose arrows route each outcome to the indexed classical bit. The grey dashed vertical bars are barriers separating the operational stages described in the text. The circuit operates as the cryptographic substrate; key-generation rates and quantum bit-error rates (QBER) are reported in Section 6.1.

Experiments were executed on the Qiskit Aer statevector simulator (with SQUANCH used to emulate the QSDC channel behaviour). A brief set of initial circuit-execution sanity checks were run on IBM Quantum hardware on the IBM Cloud free-tier student access on an earlier phase of the work; this access expired before the full evaluation reported

in Sections 6.1 and 6.4–6.6 could be run on physical hardware, and all the results in these sections are thus simulated. Replicating the present results on current-generation IBM Quantum or other physical hardware is identified as future work. Cryptographic circuits (BB84 polarisation encoding and DL04) are modelled with channel noise that produces the QBER values reported in Section 6.1; analytical circuits (QSVM, QkNN, QNN) are evaluated under ideal simulator conditions in their headline configuration, with depolarising noise injected at the simulator level for the noise-resilience experiments reported in Section 6.5. Section 6.2 (Table 2) outlines the Key-rotation policy, and HNDL adversary exposure bounds where the adversary exposure per key is computed from the measured signalling rate of approximately 1.8 kbps multiplied by the rotation interval. Section 6.4 (Table 3) demonstrate the comparative classification accuracy and training time for classical, quantum, and hybrid ML models on the unified NB-IoT signalling test set. Differences across deployment environments shown in Section 6.5 (Table 4) reflect infrastructure-level execution variability—CPU contention, scheduling jitter, and container overhead in the K8s, NS-3, and Mininet environments rather than injected quantum-channel noise.

Future work also includes scaling to physical-hardware implementation and larger feature sizes. There are two assertions that are not to be made. First, it is not claiming any overall quantum benefit: the comparison in Sections 6.4–6.6 is that of a single hybrid architecture versus classical baselines in a single NB-IoT signalling task, not quantum versus classical paradigms in general. Second, the hybrid configuration demonstrates superior noise resistance in the Kubernetes/Docker environment (Table 4, in which the F1 difference between the hybrid model and the optimal classical baseline is the largest), which is an in-experiment finding, rather than a theoretical finding. We put this provisional explanation on the more dimensional decision surface that quantum kernels generate on small, structured feature sets.

4.4. Logical Architecture Flow

The logical flow of qSiM and qSiP over dual-PLMN topology is depicted in Figure 6. qSiP cryptographic protection (red and blue channels denote quantum and classical channels respectively) is used by the N32-c handshake message and N32-f forwarded signalling message carried by the IPX between vSEPP and hSEPP. Userplane GTP-U tunnels are used for home-routed traffic along N9 and local-breakout traffic along N6.




Figure 6. Logical architecture flow showing the dual-PLMN topology, the qSiP cryptographic layer at SEPP boundaries, and the classical/quantum channels between vIPX and hIPX.

4.5. Service Domain Coverage

The framework serves the four 3GPP service domains which are enhanced Mobile Broad-band (eMBB) for sustained throughput; Ultra-Reliable Low-Latency Communications (uRLLC) for sub-10 ms response; massive Machine-Type Communications (mMTC) for dense IoT signalling; and the Billing and Charging Element (BCE), which is treated as a fourth domain because the integrity profile of charging records demands distinctive fraud-detection logic.

5. Experimental Testbed and Evaluation Setup

5.1. Environment Architecture

The evaluation environment as per Figure 7 comprises a layered virtualisation stack. Oracle VirtualBox 7.1.4 hosts Ubuntu 24.04 LTS, within which simulation (NS-3 v3.44, MATLAB R2024b/Simulink, Mininet with OVS OpenFlow 1.3), orchestration (MicroK8s v1.34, Docker v28.5.1, Vagrant v2.4.9), and 5G core (Open5GS v2.7.5 dual-PLMN) components are co-located. A lightweight C++ vC++20/Python v3.2 bridge (qiskit-interface.cc/.py) serialises NS-3 protocol parameters to JSON, invokes external Qiskit/SQUANCH scripts to generate and execute QASM circuits on IBM Quantum simulators, and injects measurement outcomes back into NS-3's event loop preserving strict synchronisation between quantum and classical event flows.




Figure 7. Testbed environment architecture: VirtualBox host, MicroK8s pods (Quantum Engine, AI/ML Classifier, SEPP/IPX Gateway, SDN Controller, STRIDE-L Threat Modelling), NS-3, MATLAB/Simulink, Mininet, and Open5GS HPMN/VPMN.

All quantum components of the experimental environment are software-simulated (Qiskit Aer for circuit simulation; SQUANCH for quantum-channel emulation); no runs were performed on physical IBM Quantum or other quantum hardware. The classical components of the environment (NS-3, Mininet, Open5GS, MicroK8s, Vagrant) execute as real software within the virtualised testbed.

5.2. Dual-PLMN Open5GS Testbed

The dual-PLMN deployment topology in Figure 8, consists of two different 5G PLMNs—home PLMN (HPMN, MCC:001/MNC:001) and visited PLMN (VPMN, MCC:999/MNC:070)—connected through an IPX fabric. Every PLMN has its own 5G core network functions (AUSF, UDM, UDR backed by MongoDB, SMF, AMF, PCF, BSF, NRF, NSSF, UPF) on its own FQDNs and IP addresses with an administrative subnet of 10.33.33.0/24. The qSiP cryptographic layer is provided by SEPP entities (sepp1.localdomain in the HPMN; sepp2.localdomain in the VPMN) and the northbound TLS secured signalling exchanges between PLMNs are handled by SEPP entities. PacketRusher is used for generating UE and gNB traffic, representing NB-IoT mobility between HPMN and VPMN. Home-routed (HR) scenarios are exercised as well as local-breakout (LBO) scenarios. Figure 7 shows the dual-PLMN deployment. The Kubernetes deployment testbed is an extension of an open-source, Open5GS-based roaming reference deployment [29] which is complemented by the following additions for the current study: an IPX fabric between the two PLMNs, the inclusion of quantum channel emulation using Qiskit and SQUANCH and the introduction of PacketRusher-generated NB-IoT mobility profiles at the inter-PLMN boundary. Both local-breakout (LBO) and home-routed (HR) scenarios are implemented. Performance metrics are captured via Wireshark/tshark, with telemetry (RSSI, SINR, QBER, entanglement fidelity) streamed to a unified analytics dashboard. The testbed used throughout this paper, with all topology definitions, Helm charts, MicroK8s manifests and the IP addressing scheme used in this paper, is released as an open repository [30] under the controlled-access policy described in the Data Availability Statement.




Figure 8. qSiM/qSiP dual-PLMN deployment topology. HPMN (MCC:001/MNC:001) and VPMN (MCC:999/MNC:070) run complete 5G Core Network Functions interconnected via IPX (vIPX ↔ hIPX). SEPP instances with embedded qSiP cryptographic layer terminate N32-c and N32-f. PacketRusher provides UE/gNB emulation in the visited network. Both HR (via N9) and LBO traffic scenarios are exercised.

5.3. Evaluation Protocol

The evaluation protocol is based on the FEDS evaluation framework [31] and uses one configuration matrix for three different environments (NS-3 v3.44, Mininet, MicroK8s), four roaming scenarios (HR, HR-MITM, LBO, LBO-MITM), four service domains (eMBB, uRLLC, mMTC, BCE), and seven model configurations (three classical: SVM, KNN, NN, three quantum: QSVM, QkNN, QNN, one hybrid ensemble: SVM/KNN, SVM/QNN, QSVM/NN, QSVM/KNN). All experiments are replicated using various seed values and configuration files are saved in Git-tracked Docker images so that they can be reused.

5.4. Stride-L Threat Modelling

The scenario generation is based on the taxonomy STRIDE-L (Spoofing, Tampering, Repudiation, Information Disclosure, Denial of Service, Elevation of Privilege, Leakage) [32]. For every STRIDE-L category, a man-in-the-middle attack is added to the testbed and checked at the packet level with Wireshark and tshark with the help of ASCII-based markers, such as IPX_SIGNALING, QKD-MITM, GTPU_EAVESDROP, QSDC-BEGIN, that should be detected by the artefact's detection logic.

6. Results

The Results section reports cryptographic-security findings first (Sections 6.1–6.3) and then the anomaly-detection findings (Sections 6.4–6.6). This ordering reflects the underlying research design: the cryptographic layer is the substrate on which the analytical layer operates, and the validity of the analytical layer's results depends on the integrity of that substrate.

6.1. Cryptographic Security at OSI L1/L2

On both HR and LBO roaming paths, the qSiP layer sustains a measured signalling throughput of approximately 400 Mbps and a BB84 key generation rate of roughly 1.4–1.5 kbps—sufficient to refresh per-session keys at intervals shorter than the typical N32 handshake period. Under nominal conditions, qSiP operates well within the BB84 decoy-state threshold of 3.5% (QBER of 1.8–2.1%). The total qSiP overhead per signalling exchange is 3.8–4.0 ms under nominal conditions, comfortably below the 5 ms budget that 5G timing constraints impose [33]. Under simulated MITM stress, the QBER rises to 5.9–6.4%—exceeding the decoy-state threshold and triggering the key-rotation policy specified in Section 6.2. The system continues to operate, with key-establishment latency increasing modestly to 4.4–4.7 ms but remaining within budget. Notably, the DL04 integrity-check overhead remains constant at 175–180 μ s/frame regardless of attack conditions, confirming that the per-frame integrity check is a low-latency, fixed-cost operation that does not contribute to the overhead growth observed under attack. The reduction in key generation rate under MITM (from 1.42 to 1.51 kbps to 0.98–1.02 kbps) arises from two factors: (a) the larger fraction of raw key material discarded during privacy amplification when the observed QBER approaches the decoy-state threshold, and (b) the brief renegotiation overhead each time a rotation event is triggered.

6.2. HNDL Key-Rotation Analysis

The HNDL adversary is assumed to be able to (I) record (in real time) the encrypted signalling traffic, and later (II) decrypt it at some future time when sufficient quantum capability becomes available. Two questions arise: first, can the cryptographic substrate determine that an ongoing interception is taking place; and second, what proportion of the traffic is still usable for future quantum decryption. For the first question, BB84's information-theoretic eavesdropping detection is used in qSiP through continuous mon-

itoring of the QBER (Section 6.1). The second is addressed by the key-rotation policy described here.

If the observed QBER on a quantum channel exceeds the decoy-state threshold, if the SDN controller receives a session anomaly reported by the response module, or if a specified time period since the last rotation has elapsed, an explicit key-rotation event is triggered. The default interval was set to 30 s because: (I) it explicitly limits the ‘harvested’ data volume available to an HNDL adversary to a single key’s worth of traffic (~5.4 KB at the measured signalling rate of ~1.8 kbps), so that even after future quantum decryption no more than this bounded fragment is recoverable per recovered key; (II) it minimises operational overhead (two rotation events per minute, or ~3% of signalling load, with each measured rotation completing in 1.42 s on average—mean over $n = 30$ runs, $\sigma = 0.18$ s—which is well below the 30 s interval); and (III) it is well within the minimum N32 session lifetime: TLS handshake optimisations have demonstrated sub-second renegotiation, while default full handshakes operate comfortably at 30 s granularity [22]. Under nominal conditions, each ephemeral key therefore protects a window of ~30 s of signalling traffic; under MITM-triggered rotation, the protected window is shorter. Even with future quantum capability, the adversary recovers only the traffic associated with one ephemeral key per recovery: the next rotation interval requires fresh key recovery, and the rotated keys are not derivable from each other. The resulting bound on adversary exposure under three rotation policies is reported in Table 2.

Table 2. Key-rotation policy and HNDL adversary exposure bounds. Adversary exposure per key is computed from the measured signalling rate of approximately 1.8 kbps multiplied by the rotation interval.

Rotation Policy	Rotation Interval (s)	Adversary Exposure per Key (KB)	Operational Overhead (Rotation Events/Minute)
Conservative (high security)	10	≈1.8	6
Default (balanced)	30	≈5.4	2
Permissive (low overhead)	60	≈10.8	1
MITM-triggered (event-driven)	variable (mean ≈ 12)	≈2.2 (mean)	≈5

Bold marks the best-performing value per metric.

Under the default (balanced) policy, the HNDL adversary can only recover about 5.4 KB of signalling traffic for each key recovered and each rotation results in a newly generated key that is not related to previous keys. Together with the eavesdropping-detection guarantee of BB84, this restricts the maximum HNDL exposure over the course of an operational period in the future to one rotation interval. The result expands previous work on the HNDL analysis reported in previous qSiM/qSiP research, which did not include an explicit analysis of the rotation-policy scores.

6.3. N32 Semantic Preservation

The signalling format using 3GPP JOSE-protected JSON is specified in 3GPP TS 29.573 [2] and the security architecture (TLS-protected N32-c handshake; PRINS or JWE-protected forwarded messages) is specified in 3GPP TS 33.501 [1]. The cryptographic substrate needs to maintain the semantics of the message format and the handshake throughout the session (end-to-end) to ensure that legacy SEPP implementations on the partner PLMN can interoperate with qSiP as is.

The trace files (qsim_roaming.pcap) were captured using Wireshark and compared. Three sets of captures were taken: a TLS-only trace (baseline); a qSiP-protected trace (no MITM); a qSiP-protected trace (simulated MITM). A comparison of the N32-c handshake messages (sec_param_exchange, JWS-signed identifiers, mutual authentication) and

N32-f forwarded messages (JOSE headers, JWE-encrypted payloads, IEs) showed that all messages in each class maintained their protocol state machine ordering and were retained in their same byte-level format. In the qSiP layer, a quantum-protected envelope is added to the encapsulated 3GPP message format. The signatures injected by STRIDE-L (IPX_SIGNALING, QKD-MITM, GTPU_EAVESDROP, QSDC-BEGIN) were captured and classified as expected, with a mean of 1.42s (median 1.39s, $\sigma = 0.18$ s, IQR 1.29–1.54 s, range 1.08–1.81 s, $n = 30$ independent runs, 95% CI for mean 1.36–1.49 s) after being injected in the qSiP-protected MITM case.

6.4. Comparative Machine Learning Performance

So far, we have shown that the cryptographic substrate is in good working order (Sections 6.1–6.3) and now we report the results of the anomaly detection. The per-model performance on the unified NB-IoT signalling test set is given in Table 3. It is important to note that quantum ML classifiers discussed in this section detect attacks, rather than encrypting signalling. The qSiP cryptographic layer reported in Section 6.1 encrypts the data.

Table 3. Comparative classification accuracy and training time for classical, quantum, and hybrid ML models on the unified NB-IoT signalling test set. Hybrid configurations combine the corresponding classical and quantum pair (e.g., hSVM = SVM + QSVM).

Model	Type	Accuracy (%)	Training Time (s)
SVM	Classical	91.3	5.2
KNN	Classical	89.4	2.1
NN	Classical	92.5	6.7
QSVM	Quantum	95.7	9.3
QkNN	Quantum	93.2	8.1
QNN	Quantum	94.6	10.5
hSVM (SVM + QSVM)	Hybrid	98.2	10.8
hKNN (KNN + QkNN)	Hybrid	97.5	9.9
hNN (NN + QNN)	Hybrid	98.7	11.5
Hybrid Ensemble	Hybrid Avg.	98.8	11.2

Bold marks the Model.

In Table 3 the classical baselines achieve 89.4–92.5% accuracy with training times between 2.1 s (KNN) and 6.7 s (NN). Quantum-enhanced models improve accuracy to 93.2–95.7% at the cost of training times increasing to 8.1–10.5 s—a trade-off consistent with known constraints of variational quantum circuit optimisation. Hybrid pairings deliver the most persuasive results: hNN (NN + QNN) achieves peak accuracy of 98.7% in 11.5 s, while hSVM (SVM + QSVM) offers the best accuracy-to-efficiency balance at 98.2% in 10.8 s. The hybrid ensemble sustains 98.8% mean accuracy in 11.2 s—a 6.3-percentage-point improvement over the best classical baseline (NN at 92.5%) and 3.1 points over the best quantum-only model (QSVM at 95.7%). The hybrid ensemble combines the three classical–quantum pairs (hSVM, hKNN, hNN) using soft voting: each constituent classifier outputs a class-probability vector $P(y|x)$, the three vectors are averaged element-wise with equal weight ($w_i = 1/3$), and the final label is the argmax of the averaged probabilities. We chose soft voting over hard voting (majority class) because it preserves the confidence margins of the underlying models, and over stacking because it avoids introducing an additional meta-learner that would obscure the contribution of the quantum components to the ensemble output. This is consistent with ensemble-learning theory: the hybrid combines complementary inductive biases—classical models’ fast, low-variance kernel decisions with quantum models’ capacity to capture non-linear high-dimensional structure—yielding a strictly dominant Pareto frontier. All ML training and inference, and

all Qiskit Aer statevector simulations, were executed on a single workstation running Ubuntu 24.04 inside Oracle VirtualBox 7.1.4: an Intel Core i7-13700H CPU (14 cores, 20 threads) with 32 GB RAM and no GPU acceleration. Qiskit Aer was version 0.14.x, SQUANCH the latest community release, scikit-learn 1.4.x for the classical baselines, and PennyLane 0.36.x for the variational QNN. The training-time figures in Table 3 reflect this single-CPU configuration; deployments with GPU acceleration would reduce the training-time gap between classical and quantum-enhanced models, though our focus here is on relative model behaviour rather than absolute throughput.

6.5. Cross-Environment Classification Performance

Table 4 reports classification performance across three deployment environments (NS-3 v3.44, Mininet, MicroK8s/Docker). Four metrics are reported per row: accuracy, precision, recall, and F1-Score (the four columns to the right of the model name), allowing each environment to be assessed on the same set of standard classification metrics rather than accuracy alone. The results demonstrate that the hybrid configuration is the best or joint-best in every environment, and that quantum-only configurations (QSVM, QkNN, QNN) outperform their classical counterparts in every environment but at a higher training-time cost.

Table 4. Cross-environment classification performance across NS-3, Mininet, and MicroK8s. Quantum and hybrid configurations consistently outperform classical baselines in every environment.

Environment	Model	Accuracy	Precision	Recall	F1-Score
Mininet	NN (C)	0.960	0.960	0.960	0.960
Mininet	KNN (C)	0.940	0.940	0.940	0.940
Mininet	SVM (C)	0.950	0.950	0.950	0.950
Mininet	QNN (Q)	0.980	0.980	0.980	0.980
Mininet	QkNN (Q)	0.990	0.990	0.990	0.990
Mininet	QSVM (Q)	0.995	0.995	0.995	0.995
Mininet	Hybrid	0.985	0.985	0.985	0.985
NS-3	NN (C)	0.970	0.970	0.970	0.970
NS-3	KNN (C)	0.950	0.950	0.950	0.950
NS-3	SVM (C)	0.960	0.960	0.960	0.960
NS-3	QNN (Q)	0.990	0.990	0.990	0.990
NS-3	QkNN (Q)	0.995	0.995	0.995	0.995
NS-3	QSVM (Q)	0.999	0.999	0.999	0.999
NS-3	Hybrid	0.995	0.995	0.995	0.995
K8s	NN (C)	0.950	0.950	0.950	0.950
K8s	KNN (C)	0.900	0.900	0.900	0.900
K8s	SVM (C)	0.930	0.930	0.930	0.930
K8s	QNN (Q)	0.960	0.960	0.960	0.960
K8s	QkNN (Q)	0.980	0.980	0.980	0.980
K8s	QSVM (Q)	0.990	0.990	0.990	0.990
K8s	Hybrid	0.980	0.980	0.980	0.980

Bold marks the Environment.

Three observations have emerged based on Table 4. First, NS-3 consistently yielded the highest accuracy across all model families, reflecting its more controlled packet-level simulation environment. QSVM in NS-3 gears towards near-perfect classification at 0.999 F1. Secondly, Kubernetes orchestration introduces measurable noise. Classical KNN drops 4–5 percentage points (from 0.950 in NS-3 to 0.900 in K8s) reflecting its sensitivity to container-induced jitter and the limits of Euclidean distance under non-stationary feature distributions. Third, quantum models are substantially more resilient to this noise: QSVM degrades only 0.009 (0.999 → 0.990) between NS-3 and Kubernetes and QkNN by 0.015. This resilience is attributable to the physics-informed feature embeddings implicit in quan-

tum kernel methods, which preserve discriminative geometry under noise perturbations that degrade classical distance metrics.

6.6. Consolidated Cross-Environment KPI Analysis

Figure 9 shows the consolidated KPI panel comparing classification accuracy, F1-Score, latency, and packet loss across the three environments. The hybrid model maintains the highest accuracy and F1 across all three environments, with latency overhead within the 5 ms qSiP budget reported in Section 6.1.




Figure 9. Consolidated KPI view across the three deployment environments. Hybrid sustains 98.5–99.5% accuracy with latency comparable to or better than quantum-only configurations.

Several quantitative observations are based on a direct extension of Figure 9. Latency: Classical ML attains latencies of 11–14 ms across environments, quantum attains 9–11 ms, and hybrid attains 7.5–9 ms—a sustained 3–5 ms reduction across environments that is operationally decisive for uRLLC compliance. Packet loss: Classical is 2.5–4.0%, quantum is 1.2–2.0%, and hybrid is 0.4–0.8%—an order of magnitude better, with a significant impact on session reliability. Throughput: Classical is 175–185 packets/s, quantum is 185–195 packets/s, and hybrid is 200–205 packets/s. Accuracy: the almost uniform 94–99% band visually compresses what are statistically strong differences, as confirmed below.

The Kubernetes environment also requires certain consideration. Classical ML performs worse than all other K8s on all metrics (maximum latency 14 ms, maximum packet loss 4%) and represents noise due to containerisation. The quantum and hybrid models are significantly more robust as the hybrid model continues to reach 9 ms latency, 0.8% loss and over 98% accuracy. This has direct consequences on the 5G deployments of cloud-native applications as operators move core functionalities into containerised applications, and the classical ML components typically deployed to detect signalling anomalies will experience accuracy decay that hybrid ML will clearly overcome.

7. Statistical Analysis

7.1. Analysis of Variance (ANOVA)

One-way ANOVA was applied across the three configurations (classical, quantum, hybrid) on each of the three deployment environments. The hybrid > {classical, quantum} contrast was statistically significant at $p < 0.05$ in all three environments: Mininet $F(2,6) = 19.27$, $p = 2.45 \times 10^{-3}$; NS-3 $F(2,6) = 20.62$, $p = 2.05 \times 10^{-3}$; K8s $F(2,6) = 8.39$, $p = 1.83 \times 10^{-2}$. ANOVA was further applied across the four 3GPP service domains (eMBB, uRLLC, mMTC, BCE), returning $F(3,8)$ values from 41.68 (eMBB) to 73.32 (BCE) with $p < 0.001$ in every domain—indicating that the hybrid advantage holds across both deployment environments and service domain workloads.

The results of the ANOVA in Figure 10 show that model type can have a very significant influence on accuracy in each slice of the experiment. These magnitudes of effects are especially high in the BCE and mMTC domains ($F = 73.32$ and 66.44 respectively), where it appears that the hybrid advantage is most meaningful when either billing-record integrity or dense-device signalling is the issue at hand and this is exactly the situation in which it has the most operational value to IPX operators.




Figure 10. ANOVA means by service domain (eMBB, uRLLC, mMTC, BCE). Hybrid uniformly outperforms classical and quantum across all four domains.

The Figure 11, largest F-statistic (8.39) occurs in the Kubernetes environment, but with an interesting level of significance ($\alpha = 0.05$) of 1.83×10^{-2} , indicating a higher level of inherent variation in containerised execution, not the weakening of the underlying architectural hierarchy.




Figure 11. ANOVA means by deployment environment (Mininet, NS-3, K8s). Hybrid sustains its advantage in every environment.

7.2. Wilcoxon Signed-Rank

The non-normal distribution of the per-run results led to the use of a Wilcoxon signed-rank test for both the hybrid-vs-classical and hybrid-vs-quantum contrasts for

the four anomaly-detection metrics (accuracy, precision, recall, F1). All eight contrasts were significant ($p < 0.05$), confirming that the difference between the hybrid configuration is not due to distributional skew. Signed-rank paired Wilcoxon tests between each hybrid family to the corresponding classical counterpart (hSVM vs. SVM, hKNN vs KNN, hNN vs. NN) proved significant improvements in all the environments and domains ($p < 0.05$). Hybrid vs. quantum analogues (e.g., hSVM vs. QSVM) pairwise comparison revealed comparable gains in the hybrid model which although small in magnitude were significant in most panels. Figure 12 presents the average performance across environments, grouped by type (classical, quantum, hybrid). Hybrid ML outperforms both classical and quantum ML consistently across all metrics, with statistical significance confirmed by ANOVA and Wilcoxon tests.




Figure 12. Per-classifier detection accuracy by service domain (eMBB, uRLLC, mMTC, BCE—Billing and Charging Evolution), with paired Wilcoxon significance. Within each panel the grouped bars give the accuracy of the Classical, Quantum, and Hybrid configurations for the SVM, KNN, and NN classifiers, and the connecting lines trace the step from Classical to Quantum and from Quantum to Hybrid. Each panel title reports the one-sided paired Wilcoxon signed-rank p -values for the hybrid-vs-classical, $p(H>C)$, and hybrid-vs-quantum, $p(H>Q)$, contrasts, evaluated independently on each domain; smaller values indicate a more consistent hybrid advantage.

Figure 13 presents the hybrid versus baseline comparison resolved by deployment environment.




Figure 13. Wilcoxon by deployment environment (Mininet, NS-3, K8s). The vertical axis reports the Wilcoxon signed-rank test statistic for the per-run hybrid vs classical and hybrid vs quantum contrasts (larger positive values indicate a stronger hybrid advantage), evaluated independently on each environment.

Figure 14 aggregates the results of the preceding two figures into a single summary view: for each of the four metrics (accuracy, precision, recall, and F1 score), the bars show

the mean score of the Classical, Quantum, and Hybrid configurations averaged over the three deployment environments (NS 3, Mininet, and K8s).




Figure 14. Average ML performance metrics across all environments (NS-3, Mininet, K8s). Colour conventions: classical models are in orange, quantum models are in blue, and hybrid models are in green. The chart aggregates per-environment means for each of the three configurations.

7.3. Verdict on H_0 and H_1

The cryptographic-security results (Sections 6.1–6.3) and the anomaly-detection results (Sections 6.4–6.6), taken together with the statistical tests reported above, support rejection of the null hypothesis H_0 in favour of the alternative H_1 : the integration of QKD/QSDC at L1/L2 with hybrid quantum–classical ML classification measurably improves the security and operational performance of 5G IRN signalling compared to classical-only baselines.

In Figure 15, top scores in interface-level security were on the user-plane N9 channel (≈ 0.93) and the control-plane N32 interface (≈ 0.88) indicating the direct positive impact of the SEPP-layer embedding of qSiP. The QKD and QSDC mechanisms had a score of around 0.80 and 0.60 respectively with the latter being lower because of fidelity sensitivity of entanglement-based teleportation to channel noise. The fact that HNDL adversaries still obtain ciphertext under quantum-secure configuration despite the score being about 0.40 could be considered as an indication that the HNDL adversary is still to take possession of the ciphertext, and could hold on to the ciphertext indefinitely, with the low score therefore being controlled by a temporal time limit between the key-rotation events rather than the failure of the cryptographic primitive. The attack impact levels that indicate the extent to which an active attacker impacts normal service operation were low in all the injected scenarios: N9 MITM = 0.08, N32 MITM = 0.15, quantum-eavesdropping = 0.20, and HNDL = 0.15. These impact numbers attest to the fact that the classification pipeline of qSiM informs the adversarial activity prior to significant service degradation. Roaming-path latencies illustrate the security–performance trade-off. Even when using MITM the home-routed paths do not surpass 80 ms, showing that the key-exchange overhead with qSiP is not catastrophically accumulated with quantum-secure processing. Local-breakout baselines are around 30 ms and go to around 40 ms with MITM a 33% overhead in latency due to extra verification of anomaly classifications. This overhead is operationally tolerable in the case of NB-IoT workloads. The pie chart of the harvested-data composition confirms the approximate distribution between N32 control-plane messages and GTP-U user-plane data, reflecting the separation between intercepted traffic between the control plane (protected) and the user-plane (protected) domains as is typical of N32 hardening proposals.




Figure 15. Consolidated security and billing-fraud dashboard. (**Top-left**): security metric scores (per-interface scores for N9, N32, QKD, QSDC, and HNDL bounded-exposure). (**Top-right**): attack impact levels (N9 MITM, N32 MITM, quantum-eavesdropping, HNDL). (**Bottom-left**): roaming-path latencies (HR baseline, HR with MITM, LBO baseline, LBO with MITM). (**Bottom-right**): harvested-data types under simulated attack—the two pie segments are labelled “N32 Control Plane (Protected)” (the larger share, dominated by N32-c handshakes and N32-f forwarded signalling) and “GTP-U User Plane (Protected)” (the remaining share, comprising GTP-U user-plane data on the N9 interface for HR scenarios). Both segments correspond to traffic protected by the qSiP layer; the chart shows the relative composition of what an HNDL adversary would harvest, not what is in the clear.

The most educative outcome in Figure 15 is the billing-fraud detection. When using qSiM in the standard roaming (HR and LBO without active attack) configuration, the classification pipeline detects the most fraudulent billing messages of the home-routed and local-breakout scenarios respectively of 11 out of 15 and 10 out of 15. Detection rates decline significantly under active MITM (to around 6–7 of 15 injected fraud messages) and point to the fact that advanced attackers with the capability to modify classification features during transit can suppress some fraction of the fraud cues. It is a substantive result as it demonstrates that the fraud-detection capability of qSiM is a fact but not free and that content-layer ML classification (detecting fraud that avoids the transport defences) and defence-in-depth combining transport-layer quantum security (preventing the MITM in the first place) are architecturally necessitated. Precisely this combination of protection accounts to the N32 and N9 security scores of 0.88 and 0.93 respectively.

Mechanistic analysis of MITM-induced detection degradation. The substantial fall in billing-fraud detection under active MITM warrants a causal explanation. The hybrid classifier was trained on NB-IoT features—packet size, inter-arrival time, source/destination pair, and modulation class under nominal and passive-adversary conditions, where injected fraud messages carry distinctive signatures (anomalous inter-arrival cadence and enlarged JOSE payloads). Under active MITM the adversary manipulates packets in transit before feature extraction: payload fields are silently modified to fall within the normal size range, and inter-arrival timing is perturbed by relay buffering, smearing the distribution toward

the legitimate mean. The two most discriminative feature dimensions are thereby actively normalised, rendering injected fraud statistically indistinguishable from legitimate traffic, a targeted feature-space evasion attack rather than mere added noise.

The quantum kernel components (QSVM, QkNN) are more resilient because the ZZ feature map embeds the inputs into a higher-dimensional Hilbert space in which the adversary's normalisation of the original four features does not necessarily collapse separability; the effect is partial, which accounts for the hybrid configuration retaining roughly six to seven detections of fifteen injected fraud messages versus four to five for the purely classical baseline under MITM. Robustness to this attack therefore requires adversarially trained classifiers using MITM-perturbed feature distributions, or features that the relay cannot silently normalise such as sequence-number entropy, TLS record-layer timing jitter, or the QBER time-series from the quantum sub-channel, both of which are identified as future work in Section 9.

8. Discussion

8.1. The Cryptographic/Analytical Distinction

The cryptographic and the analytical aspects of quantum techniques are clearly distinguished in this paper. The cryptographic role is assigned to qSiP: BB84 QKD is used for key establishment that is detectable by eavesdropper and DL04 QSDC is used for direct secure communication at L2. No anomaly classification is carried out by either. The analytical function is a hybrid ML classifier that inputs the signalling features and classifies them as a normal, abnormal, or fraudulent signal. It does not encrypt anything. Both roles are complementary, the framework requires the use of both, and conclusions about the behaviour of the artefact should be made with respect to the role under consideration. Sections 6.1–6.3 assess the cryptographic role, and Sections 6.4–6.6 assess the analytical role.

8.2. Threats to Validity

Construct validity is invalidated by the ability of the NS-3 simulator to offer fine-grained control over NB-IoT PHY/MAC parameters, but real-world implementations may demonstrate behaviour that is not simulated. Background traffic from the Mininet hosts may also cause some variability in latency, which is not related to the quantum channel; this was controlled by keeping the load of the hosts fixed and running experiments multiple times with different seeds. The scenario studied (single operator IPX) is an impediment to the external validity; multi-operator and multi-hop SEPP topologies may bring synchronisation and scaling challenges.

8.3. Limitations

All quantum protocols have been simulated in software using the Qiskit Aer statevector simulator (Qiskit Aer 0.14.x) for circuit-level evaluation and SQUANCH for QSDC channel emulation, with PennyLane 0.36.x used for the variational QNN training loop. The characteristics of the physical QKD hardware (polarisation drift, detector inefficiencies, dark counts, after-pulsing) have been parameterised but not empirically measured on a real device. The noise in the quantum channel was modelled as Poisson-distributed phase errors injected at the simulator level; more complicated environmental effects such as temperature fluctuations, fibre-length variability, and detector saturation were not modelled. A brief earlier set of circuit-execution sanity checks was performed on real IBM Quantum hardware via the IBM Cloud free-tier student access; that access expired before the full evaluation reported in Sections 6.1 and 6.4–6.6 could be conducted on physical hardware, and all numerical results in those sections are therefore based on simulation. Replication

of the present results on current-generation IBM Quantum or other physical hardware is identified as future work.

To make the scope of the simulation explicit, we distinguish three tiers of fidelity. (a) Parameters instantiated directly in the simulator with empirically grounded values: Poisson-distributed phase errors producing QBER values consistent with published fibre-based BB84 experiments, the 3.5% decoy-state threshold, and a SQUANCH-emulated QSDC channel fidelity of 0.98. (b) Parameters abstracted away by the simulator: detector dark-count rates, after-pulsing, polarisation-mode dispersion, wavelength-dependent fibre attenuation, and photon-source spectral purity. (c) Effects entirely outside the simulation envelope: component-level thermal drift over operational timescales, multi-photon (Trojan-horse) pulse vulnerability, and the latency of real quantum-memory or optical-switch hardware. Accordingly, the numerical results in Sections 6.1 and 6.4–6.6 should be read as a simulation-based proof of concept rather than a hardware-validated measurement.

8.4. Sensitivity

Parametric sweeps in NS-3 demonstrate that classification accuracy drops by up to 5% with QBER greater than 7% and emulated latency does not exceed 10 ms with emulated simulation-bridge event-queue depth set to 64 events or greater. This parameter is a classical software property of the Qiskit/SQUANCH–NS-3 interface bridge, representing the number of outstanding serialised N32 events the bridge may queue before back-pressuring the NS-3 scheduler and not a physical quantum-state buffer; below the threshold the NS-3 event loop stalls waiting for circuit-execution results, inflating emulated latency. The 64-event figure therefore reflects the simulator’s concurrency capacity rather than any photonic-hardware constraint, and would not translate directly to a physical-hardware implementation. The sensitivities outline the operating range of the results and would need further engineering to deploy beyond the range.

8.5. Scalability

Edge-node inference of the quantum-enhanced classifier requires roughly 50 MB of model storage and a single-core CPU for real-time operation at 100 packets per second. This footprint is small enough that telecom operators could deploy the artefact in existing edge-computing clusters and phase in the quantum-security modules incrementally, minimising CAPEX.

9. Conclusions and Future Work

In this paper, a hybrid quantum–classical anomaly detection framework for 5G IRN signalling was developed and validated in simulation combining QKD and QSDC at OSI Layers 1 and 2 with a hybrid classical-plus-quantum ML classifier for anomaly detection of signalling anomalies and billing fraud. The cryptographic-security results demonstrate that qSiP can keep BB84 QBER within the decoy-state range under nominal conditions, it invokes an explicit key rotation under MITM stress and by using the rotation policy, it limits the exposure of the adversary to the key to one rotation interval. The 3GPP message-format and handshake semantics are preserved in N32 packet captures. Anomaly-detection results reveal 98.5–99.5% hybrid accuracy in three deployment environments. The hybrid detector significantly reduces undetected billing fraud compared to classical baselines (e.g., 11 out of 15 and 10 out of 15 for classical under HR), though active MITM attacks still evade detection in approximately half of injected fraud cases, highlighting the need for the combined QKD + ML defence. The evidence supports the single objective in Section 1.1 and the alternative hypothesis H_1 is accepted instead of H_0 .

There are three commitments that define the contributions. First, the cryptographic and analytical functions of quantum processes remain separated: QKD/QSDC encrypt, while quantum-enhanced ML classifiers detect. Secondly, the cryptographic results are reported first (Sections 6.1–6.3) and the analytical results second (Sections 6.4–6.6), corresponding to the layered design. Third, the HNDL adversary is addressed with explicit key-rotation evidence rather than static QKD/QSDC scores. Two scope statements should be read alongside these contributions. (I) The quantum evaluation reported in this paper was conducted in simulation (Qiskit Aer statevector simulator and SQUANCH for QSDC channel emulation) rather than on physical quantum hardware; a brief earlier set of circuit-execution sanity checks was run on real IBM Quantum hardware via the IBM Cloud free-tier student access, but that access expired before the full evaluation campaign could be executed on physical devices, so the headline numerical results in Sections 6.1 and 6.4–6.6 are simulation-based. (II) The scalability and real-world deployment constraints of qSiM/qSiP have not been fully evaluated in this paper: the cross-environment results (Mininet, NS-3, MicroK8s/Docker) demonstrate behaviour under three increasingly realistic execution substrates, but multi-operator IPX deployment, geographically distributed SEPP topologies, photon-source/detector hardware integration, and long-running production operation remain open. Future development will accordingly address (a) multi-hop SEPP topologies and multi-operator IPX peering; (b) replacement of simulated quantum components with physical QKD/QSDC links, including detector calibration and polarisation-drift compensation; (c) formal verification of the qSiP state machine against 3GPP TS 33.501; and (d) reapplication of the framework to adjacent domains, including 6G slicing and inter-satellite signalling. Future development will also address (e) adversarially robust training and feature augmentation for MITM-resilient billing-fraud detection—including training sets augmented with MITM-perturbed feature distributions and features the relay cannot silently normalise, such as sequence-number entropy, TLS record-layer timing jitter, and the QBER time-series from the quantum sub-channel.

Because the single-operator IPX scenario studied here limits external validity, we set out four concrete directions for multi-operator and multi-hop extension. First, distributed key management across federated QKD networks: in a multi-hop path (UE $\rightarrow$ vSEPP $\rightarrow$ IPX-A $\rightarrow$ IPX-B $\rightarrow$ hSEPP) each link requires its own BB84 key pair, raising the question of whether a hierarchical key-management service can coordinate ephemeral-key lifetimes without a single point of failure or correlated rotation events that leak timing side-channels, taking ITU-T Y.3803 trusted-relay models as a starting point. Second, inter-SEPP quantum channel establishment: deploying a fibre or free-space quantum channel between SEPPs in different operator facilities raises optical-interface standardisation (absent from TS 33.501), co-existence of quantum and classical photons on shared fibre, and inter-operator trust governance at the quantum layer. Third, cumulative latency under multi-hop key chaining: each hop adds 3.6–4.7 ms of key-establishment latency, so a three-hop inter-continental path could accumulate 10–15 ms, approaching the uRLLC budget, and latency-hiding strategies such as pre-generated key pools warrant study. Fourth, operational-cost modelling: a cost-benefit model comparing HNDL risk exposure with and without quantum channels, expressed through the expected monetary value of a successful HNDL decryption event, would give operators quantitative investment guidance.

Author Contributions: Conceptualisation, T.N. and B.K.; methodology, T.N.; software, simulation, and testbed, T.N.; validation, T.N. and B.K.; formal analysis, T.N.; data curation, T.N.; writing—original draft preparation, T.N.; writing—review and editing, T.N. and B.K.; supervision, B.K. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: Data and code supporting this study are available from the corresponding author on reasonable request. Owing to the dual-use nature of certain components (L1/L2 capture and injection scripts, MITM positioning code, unsanitised handshake captures), access is restricted under a responsible-use agreement. All testbed PLMNs use 3GPP-reserved test MCC/MNC values; no real operator data is included.

Acknowledgments: The author thanks the academic supervisor, industry experts consulted during problem identification, and the open-source communities behind NS-3, Open5GS, Qiskit, SQUANCH, PennyLane, and Mininet, whose tools made this research possible. This work was supported by technical resources donated by CSG International, specifically the computing hardware used for all experiments and data analysis.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

3GPP	3rd Generation Partnership Project
AF	Application Function (5G Core)
AMF	Access and Mobility Management Function
ANOVA	Analysis of Variance
AUSF	Authentication Server Function
BB84	Bennett–Brassard 1984 quantum key distribution protocol
BCE	Billing and Charging Evolution
CISA	Cybersecurity and Infrastructure Security Agency (US)
CkDT	Concept–Knowledge Design Theory
CML	Classical Machine Learning
DL04	Deng–Long 2004 quantum secure direct communication protocol
DN	Data Network
DSRM	Design Science Research Methodology
eMBB	enhanced Mobile Broadband
ENISA	European Union Agency for Cybersecurity
GTP-U	GPRS Tunnelling Protocol—User plane
HNDL	Harvest Now, Decrypt Later
HPMN/HPLMN	Home Public Mobile Network/Home Public Land Mobile Network
HR	Home-routed (roaming scenario)
hSEPP	home Security Edge Protection Proxy
IPX	IP eXchange (inter-operator carrier backbone)
IRN	Interconnect and Roaming Network
JOSE	JSON Object Signing and Encryption
JWE	JSON Web Encryption
KMS	Key Management Service
LBO	Local breakout (roaming scenario)
MITM	Man-in-the-middle
mMTC	massive Machine-Type Communication
MNO	Mobile Network Operator
N1	Reference point between UE and AMF
N2	Reference point between (R)AN and AMF
N3	Reference point between (R)AN and UPF
N32	Reference point between SEPPs (inter-PLMN signalling)
N32-c	N32 control sub-interface (handshake)
N32-f	N32 forwarded-signalling sub-interface (JOSE/JWE-protected)
N4	Reference point between SMF and UPF
N6	Reference point between UPF and Data Network

N9	Reference point between two UPFs (used for inter-PLMN user-plane tunnelling)
Naf	Service-based interface of the Application Function
Namf	Service-based interface of the AMF
Nausf	Service-based interface of the AUSF
NB-IoT	Narrowband Internet of Things
NFV	Network Function Virtualisation
NIST	National Institute of Standards and Technology (US)
Nnrf	Service-based interface of the NRF
Nnssf	Service-based interface of the NSSF
Npcf	Service-based interface of the PCF
NRF	Network Repository Function
NSA	National Security Agency (US)
Nsmf	Service-based interface of the SMF
Nsmsf	Service-based interface of the SMSF
NSSF	Network Slice Selection Function
Nudm	Service-based interface of the UDM
OFDM	Orthogonal Frequency-Division Multiplexing
PCF	Policy Control Function
PLMN	Public Land Mobile Network
PRINS	Protocol for N32-interconnect Security
QAM	Quadrature Amplitude Modulation
QBER	Quantum Bit Error Rate
QKD	Quantum Key Distribution
QkNN	Quantum K-Nearest Neighbour
QML	Quantum Machine Learning
QNN	Quantum Neural Networks
QSDC	Quantum Secure Direct Communication
qSiM	Quantum Signalling Model (this work)
qSiP	Quantum Signalling Protocol method (this work)
QSVM	Quantum Support Vector Machine
RAN	Radio Access Network
SBA	Service-Based Architecture (5G Core)
SDN	Software-Defined Networking
SEPP	Security Edge Protection Proxy
SMF	Session Management Function
SMSF	Short Message Service Function
STRIDE-L	Spoofing, Tampering, Repudiation, Information Disclosure, Denial of Service, Elevation of Privilege, Leakage (extended threat model)
UDM	Unified Data Management
UE	User Equipment
UPF	User Plane Function
URLLC	Ultra-Reliable Low-Latency Communication
VPMN/VPLMN	Visited Public Mobile Network/Visited Public Land Mobile Network
vSEPP	visited Security Edge Protection Proxy

References

1. TS 33.501; Security Architecture and Procedures for 5G System. 3rd Generation Partnership Project: Valbonne, France, 2024.
2. TS 29.573; 5G System; Public Land Mobile Network (PLMN) Interconnection. 3rd Generation Partnership Project: Valbonne, France, 2024.
3. FS.21; Interconnect Signalling Security Recommendations. GSMA Fraud and Security Group: London, UK, 2025.
4. Mosca, M. Cybersecurity in an Era with Quantum Computers: Will We Be Ready? *IEEE Secur. Priv.* **2018**, *16*, 38–41. [[CrossRef](#)]

5. Cybersecurity and Infrastructure Security Agency (CISA); National Security Agency (NSA); National Institute of Standards and Technology (NIST). *Quantum-Readiness: Migration to Post-Quantum Cryptography (Joint Cybersecurity Factsheet, TLP:CLEAR)*; CISA: Washington, DC, USA, 2023. Available online: [https://www.cisa.gov/sites/default/files/2023-08/Quantum%20Readiness_Final_CLEAR_508c%20\(3\).pdf](https://www.cisa.gov/sites/default/files/2023-08/Quantum%20Readiness_Final_CLEAR_508c%20(3).pdf) (accessed on 3 March 2024).
6. Sivanathan, A.; Gharakheili, H.H.; Loi, F.; Radford, A.; Wijenayake, C.; Vishwanath, A.; Sivaraman, V. Classifying IoT Devices in Smart Environments Using Network Traffic Characteristics. *IEEE Trans. Mob. Comput.* **2019**, *18*, 1745–1759. [CrossRef]
7. Abbas, A.M.; Youssef, K.Y.; Zekry, A.; Mahmoud, I.I. NB-IoT optimisation: Holistic view for smart cities applications with smart meters networks case study. *IET Commun.* **2021**, *15*, 112–125. [CrossRef]
8. Javeed, D.; Saeed, M.S.; Ahmad, I.; Adil, M.; Kumar, P.; Najmul Islam, A.K.M. Quantum-Empowered Federated Learning and 6G Wireless Networks for IoT Security: Concept, Challenges and Future Directions. *Future Gener. Comput. Syst.* **2024**, *160*, 577–597. [CrossRef]
9. Kundu, S.; Ghosh, S. Security Aspects of Quantum Machine Learning: Opportunities, Threats and Defenses. In *Proceedings of the Great Lakes Symposium on VLSI 2022 (GLSVLSI '22)*, Irvine, CA, USA, 6–8 June 2022; ACM: New York, NY, USA, 2022; pp. 463–468. [CrossRef]
10. IR.34; Inter-Service Provider IP Backbone Guidelines, Version 14.0. GSMA Permanent Reference Document. GSM Association: London, UK, 2020.
11. NG.113; 5GS Roaming Guidelines, Version 2.0. GSMA Permanent Reference Document. GSM Association: London, UK, 2021. Available online: <https://www.gsma.com/newsroom/wp-content/uploads//NG.113-v2.0.pdf> (accessed on 7 May 2026).
12. ENISA. *Signalling Security in Telecom SS7/Diameter/5G*; EU Agency for Cybersecurity: Heraklion, Greece, 2018.
13. Lutu, A.; Perino, D.; Bagnulo, M.; Frias-Martinez, E.; Khangosstar, J. A Characterization of the COVID-19 Pandemic Impact on a Mobile Network Operator Traffic. In *Proceedings of the ACM Internet Measurement Conference (IMC '20)*, Virtual Event, USA, 27–29 October 2020; ACM: New York, NY, USA, 2020; pp. 19–33. [CrossRef]
14. TS 23.501; System Architecture for the 5G System (5GS). 3rd Generation Partnership Project: Valbonne, France, 2024.
15. Holtmanns, S.; Rao, S.P.; Oliver, I. User Location Tracking Attacks for LTE Networks Using the Interworking Functionality. In *Proceedings of the IFIP Networking Conference*, Vienna, Austria, 17–19 May 2016.
16. Bennett, C.H.; Brassard, G. Quantum Cryptography: Public Key Distribution and Coin Tossing. In *Proceedings of the IEEE International Conference on Computers, Systems and Signal Processing*, Bangalore, India, 10–12 December 1984; pp. 175–179.
17. Hwang, W.-Y. Quantum Key Distribution with High Loss: Toward Global Secure Communication. *Phys. Rev. Lett.* **2003**, *91*, 057901. [CrossRef] [PubMed]
18. Deng, F.G.; Long, G.L. Secure Direct Communication with a Quantum One-Time Pad. *Phys. Rev. A* **2004**, *69*, 052319. [CrossRef]
19. Long, G.L.; Liu, X.S. Theoretically Efficient High-Capacity Quantum-Key-Distribution Scheme. *Phys. Rev. A* **2002**, *65*, 032302. [CrossRef]
20. Biamonte, J.; Wittek, P.; Pancotti, N.; Rebentrost, P.; Wiebe, N.; Lloyd, S. Quantum Machine Learning. *Nature* **2017**, *549*, 195–202. [CrossRef] [PubMed]
21. Havlíček, V.; Córcoles, A.D.; Temme, K.; Harrow, A.W.; Kandala, A.; Chow, J.M.; Gambetta, J.M. Supervised Learning with Quantum-Enhanced Feature Spaces. *Nature* **2019**, *567*, 209–212. [CrossRef] [PubMed]
22. Lastre, J.K.; Ko, Y.; Kwon, H.; You, I. Evaluating Transport Layer Security 1.3 Optimization Strategies for 5G Cross-Border Roaming: A Comprehensive Security and Performance Analysis. *Sensors* **2025**, *25*, 6144. [CrossRef] [PubMed]
23. PQ.05; Post-Quantum Cryptography for 5G Roaming Use Case. GSMA Permanent Reference Document. GSM Association: London, UK, 2025.
24. Cao, Y.; Zhao, Y.; Wang, Q.; Zhang, J.; Ng, S.X.; Hanzo, L. The Evolution of Quantum Key Distribution Networks: On the Road to the Qinternet. *IEEE Commun. Surv. Tutor.* **2022**, *24*, 839–894. [CrossRef]
25. Mehic, M.; Niemiec, M.; Rass, S.; Ma, J.; Peev, M.; Aguado, A.; Martin, V.; Schauer, S.; Poppe, A.; Pacher, C.; et al. Quantum Key Distribution: A Networking Perspective. *ACM Comput. Surv.* **2020**, *53*, 96. [CrossRef]
26. Hevner, A.R.; March, S.T.; Park, J.; Ram, S. Design Science in Information Systems Research. *MIS Q.* **2004**, *28*, 75–105. [CrossRef]
27. Peffers, K.; Tuunanen, T.; Rothenberger, M.A.; Chatterjee, S. A Design Science Research Methodology for Information Systems Research. *J. Manag. Inf. Syst.* **2007**, *24*, 45–77. [CrossRef]
28. Hatchuel, A.; Weil, B. C-K Design Theory: An Advanced Formulation. *Res. Eng. Des.* **2009**, *19*, 181–192.
29. Lavelliane. open5gs-roaming-microk8s: Open5GS Roaming Setup with PacketRusher Deployment to MicroK8s. GitHub Repository. Available online: <https://github.com/Lavelliane/open5gs-roaming-microk8s> (accessed on 2 May 2026).
30. Ngoben, T. open5gs-roaming-microk8s (Fork, Phdroaming Branch): Quantum-Secure Signalling Testbed Extending Open5GS Roaming with IPX, IPX-SEPP, and QKD/QSDC Layers. GitHub Repository, Commit 76430c5. Available online: <https://github.com/ngothe01-lab/open5gs-roaming-microk8s/tree/phdroaming> (accessed on 2 May 2026).
31. Venable, J.; Pries-Heje, J.; Baskerville, R. FEDS: A Framework for Evaluation in Design Science Research. *Eur. J. Inf. Syst.* **2016**, *25*, 77–89. [CrossRef]

32. Shostack, A. *Threat Modeling: Designing for Security*; Wiley: Indianapolis, IN, USA, 2014.
33. ITU-R. *Minimum Requirements Related to Technical Performance for IMT-2020 Radio Interface(s)*; Report M.2410-0; International Telecommunication Union: Geneva, Switzerland, 2017.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.