Computers, Materials & Continua

Tech Science Press

https://doi.org/10.32604/cmc.2026.084352

ARTICLE

Hybrid Quantum-Kernel and Quantum-Inspired Machine Learning for TDoS Early Warning in Critical Infrastructure

Carlos Rosa-Remedios*, Pino Caballero-Gil* and Jezabel Molina-Gil

Department of Computer Engineering and Systems, Higher School of Engineering and Technology, University of La Laguna, Tenerife, Spain

*Corresponding Authors: Carlos Rosa-Remedios. Email: crosarem@ull.edu.es; Pino Caballero-Gil. Email: pcaballe@ull.edu.es

Received: 21 April 2026; Accepted: 11 June 2026; Published: 13 August 2026

ABSTRACT: Increasing digitalization exposes critical infrastructure to sophisticated cyber threats, requiring new approaches to improving security and resilience. While classical machine learning techniques have shown promise in anomaly detection and threat mitigation, emerging quantum-inspired methods offer new opportunities to enhance detection capabilities by leveraging principles derived from quantum computing. The objective of this work is to propose a model for the early detection of Telephony Denial of Service attacks using a combination of classical algorithms and quantum computing-based techniques. Call records are embedded into a low-dimensional quantum feature space using spatial and temporal attributes, mapped through a quantum kernel in two modalities: Quantum-Inspired Kernel or Qiskit-Simulated Quantum Kernel. This approach enables the identification of subtle and complex patterns that can be difficult to capture using classical methods alone, without the need for fully fault-tolerant quantum hardware. The anomaly decision threshold is empirically calibrated, allowing the detector to balance early sensitivity against false-alarm control under the evaluated Public Safety Answering Points (PSAP) traffic conditions. The evaluation shows that the proposed framework can improve threat detection and system resilience in critical infrastructure contexts. The results support the use of kernel-based quantum-inspired representations as a tunable early-warning layer for PSAP traffic monitoring, while also showing that threshold calibration and operational context remain necessary before deployment.

KEYWORDS: Critical infrastructure; quantum computing; telephony denial of service; quantum machine learning

1 Introduction

Quantum computing offers promising capabilities for enhancing information security and cybersecurity [1,2]. As these technologies mature, integrating quantum-based solutions into critical infrastructure protection becomes increasingly viable and necessary to ensure the continuity of essential public services [3]. This study specifically addresses the digital resilience of citizen service centers, such as Public Safety Answering Points (PSAPs), which rely heavily on telephony as the primary access channel for the public.

These facilities are highly vulnerable to Telephony Denial of Service (TDoS) attacks, where a disproportionate influx of simultaneous calls can neutralize their response capacity [4,5]. Mitigating these threats requires systems capable of accurately distinguishing anomalous call flows from the standard temporal and geographic distribution of legitimate traffic [6]. Although classical machine learning algorithms, such as Gaussian Mixture Models [7], have traditionally been used for characterization of spatial and temporal data,

the increasing complexity of cyberattacks and the need to minimize response times require the exploration of advanced computing paradigms.

To address these challenges, this paper proposes a hybrid quantum-classical machine learning framework for unsupervised anomaly detection. The system is designed to identify TDoS-indicative deviations in PSAP call flows by establishing a baseline profile of normality derived from operational historical data. Subsequent traffic is evaluated against this baseline using a similarity metric defined by a kernel function [8,9]. Two variants for the kernel calculation are introduced and evaluated: a Quantum-Inspired Kernel and a Qiskit-Simulated Quantum Kernel. The robustness of the methodology is validated through the use of a large-scale real-world dataset, complemented by synthetic data modeling typical TDoS attack parameters.

The remainder of this paper is organized as follows. Section 2 introduces the study datasets, synthetic scenario construction, and quantum machine learning concepts. Section 3 describes the proposed framework and detection model. Section 4 reports the experimental results, including scenario validation, threshold sensitivity, classical baselines, and computational cost. Section 5 discusses the main findings and limitations. Section 6 concludes the paper and outlines future work.

2 Methods and Parameters

2.1 Foundational Premises

To develop a mechanism capable of identifying a hypothetical TDoS attack on calls directed to a PSAP, the proposed framework is based on the following premises:

- Data-driven approach. A collection of 11 datasets was compiled, including anonymized call records directed to a PSAP between 2011 and 2021. A different data set was generated for each year, which contained metadata such as call date, geographic location, call duration, and time spent in the queue before being answered (waiting time).
- Synthetic data generation to emulate TDoS attacks. To simulate a call flow representing a hypothetical TDoS attack, an algorithm was developed to generate a synthetic dataset exhibiting the characteristics typical of such an event. Given the impossibility of obtaining a dataset from an actual TDoS attack, several variants were also generated to validate the results.
- Hybrid architecture for anomaly detection. A hybrid architecture leveraging Quantum Machine Learning (QML) algorithms [10] is proposed to model the density of legitimate calls and discriminate statistical anomalies.

The subsequent sections detail the technical aspects of each of these foundational elements, beginning with a description of the dataset and the applied data cleansing processes. This is followed by an elaboration of the quantum computing concepts utilized in the implementation of QML algorithms.

2.2 Study Datasets

2.2.1 Features of the Datasets

The dataset used in this study consists of anonymized call records from the emergency call center in the Canary Islands, Spain. A total of 11 annual datasets are available, corresponding to all calls received between 2011 and 2021 in the province of Las Palmas. The 2019 dataset was initially selected as a reference to characterize normal call flow behavior, based on criteria validated by PSAP operational staff and because it was not associated with serious emergency incidents that could distort the normal profile. To ensure a robust analysis and avoid relying on a single calibration year, we later expanded the baseline to include 2018, 2021, and a combined three-year average. Since real-world TDoS traces are not publicly available, we tested our approach using several synthetic scenarios. This included new "zero-day" variants, all specifically tailored to and validated against the Canary Islands' operational environment.

The initial features of the working datasets, along with a brief description of their meaning, are detailed in Table 1.

Table 1: Definition and description of the features within the TDoS detection framework.

Column Name	Description
DATE	Timestamp of the recorded call (Year/Month/Day/Hour).
POSITION	Spatial coordinates including approximate longitude and latitude.
PHONE	Unique identification number of the calling party (anonymized).
TALK_TIME	Total duration of the call communication phase.
QUEUE_TIME	Waiting time elapsed before the call was answered.
TYPE	Categorization of the positioning data source/methodology.

2.2.2 Exploratory Data Analysis

Prior to the main analysis, an exploratory evaluation was conducted to assess overall data quality and identify potential anomalies. Table 2 summarizes these findings, detailing the size of each annual dataset (Records), feature counts (Feat.), duplicate entries, and the percentage of missing, malformed, or non-parseable coordinate fields after position parsing (Miss./Non-parseable Pos. (%)). Additionally, the sample skewness of latitude coordinates was calculated to highlight asymmetry and detect potential spatial outliers.

Table 2: Longitudinal data integrity (2011–2021).

Year	Records	Feat.	Miss./Non-Parseable Pos. (%)	Duplicates	Spatial Skew
2011	814,078	6	0.000000%	0	−9.841795
2012	662,384	6	0.000000%	2	−26.095271
2013	822,286	6	0.000000%	2	−27.882670
2014	765,836	6	0.000000%	9	6.023686
2015	759,355	6	0.000132%	20	8.977573
2016	695,905	6	0.000000%	4	2.115057
2017	605,541	6	0.000000%	3	14.075310
2018	634,097	6	0.037691%	0	8.328763
2019	631,177	6	0.000000%	0	31.305737
2020	552,082	6	0.000000%	94	32.571979
2021	725,496	6	3.572728%	5	1.218430

Based on this data-quality audit, the following observations can be drawn:

- The datasets maintain dimensional consistency throughout the complete observation period, with six features available in every yearly file (Feat. = 6).
- Positioning information is available for the vast majority of records. The audit, however, identifies a small proportion of missing, malformed, or non-parseable spatial fields in specific years. This proportion is negligible in 2015 (0.000132%) and remains low in 2018 (0.037691%), whereas 2021 presents a higher rate of unusable positioning records (3.572728%). Consequently, geolocation remains a suitable primary attribute for the proposed model, but only after explicit coordinate parsing, validation, and exclusion of invalid spatial entries.

- The duplication rate is statistically insignificant in all yearly datasets. The highest value is observed in 2020, with 94 duplicate records out of 552,082, corresponding to less than 0.02% of the dataset.

The Spatial Skew values show substantial interannual variability, indicating that the spatial distribution of calls is not symmetric and that the location of extreme observations changes across years. High positive values, such as those observed in 2019 and 2020, indicate the presence of distant records at higher latitudes relative to the annual mean. Conversely, strongly negative values, such as those observed in 2012 and 2013, reflect extreme observations below the mean latitude.

This variation in both magnitude and sign suggests that spatial outliers are not governed by a stable yearly pattern. Rather, they appear as irregular effects associated with the operational and geolocation characteristics of the PSAP data. This behavior reinforces the need for robust preprocessing and supports the use of non-linear kernel-based methods, since a classical linear model would be poorly suited to capture such changing spatial asymmetries without either overfitting to outliers or incorrectly treating legitimate spatial extremes as anomalous events.

To avoid dependence on a single normality profile, the baseline analysis includes 2018, 2019, 2021, and a combined multiyear baseline. Although 2019 was originally selected—validated by PSAP staff and free of major emergency anomalies—data from 2018 and 2021 were incorporated to address potential baseline dependence. Consequently, these three years are evaluated both as distinct single-year normality profiles and as a combined multi-year baseline.

For each candidate year, geographical coordinates were extracted from the `POSITION` feature, using the call hour as the temporal marker. This facilitates a direct comparison of spatial concentration and circadian distributions across the selected baselines. Descriptive statistics are provided in Table 3, with corresponding geographical density maps and temporal trends illustrated in Fig. 1.

Table 3: Baseline descriptive statistics for 2018, 2019, and 2021.

Feature	Mean	Std	Min	Max	Skewness	Kurtosis	n_ valid	Year	records_ original	records_ valid	records_used_ for_density_map
Latitude (Pos.)	28.228	0.397	27.736	58.059	8.329	555.296	633,858	2018	634,097	633,858	633,847
Longitude (Pos.)	−15.008	0.804	−42.387	−13.430	0.220	22.618	633,858	2018	634,097	633,858	633,847
Hour of Day	14.224	5.831	0.000	24.000	−0.508	−0.352	633,858	2018	634,097	633,858	633,847
Latitude (Pos.)	28.213	0.708	27.646	58.059	31.306	1311.422	631,177	2019	631,177	631,177	630,915
Longitude (Pos.)	−15.106	0.960	−42.387	−13.430	−9.016	269.453	631,177	2019	631,177	631,177	630,915
Hour of Day	14.148	5.883	0.000	24.000	−0.481	−0.400	631,177	2019	631,177	631,177	630,915
Latitude (Pos.)	28.181	0.348	1.616	52.534	1.218	84.027	699,576	2021	725,496	699,576	699,570
Longitude (Pos.)	−15.143	0.746	−18.098	13.274	1.179	3.552	699,576	2021	725,496	699,576	699,570
Hour of Day	14.314	5.858	0.000	24.000	−0.499	−0.335	699,576	2021	725,496	699,576	699,570

Note: *records_valid* denotes records with parseable longitude, latitude, and temporal information. *records_used_for_density_map* denotes the subset retained for the descriptive geographical density maps after excluding extreme coordinates outside the Canary Islands display domain. This additional spatial filter was used only for exploratory visualization and does not correspond to duplicate removal or to missing-position filtering. Thus, for 2019, the difference between 631,177 valid records and 630,915 records used in the density map corresponds to 262 spatial outliers outside the display domain, not to missing coordinates or duplicated records.

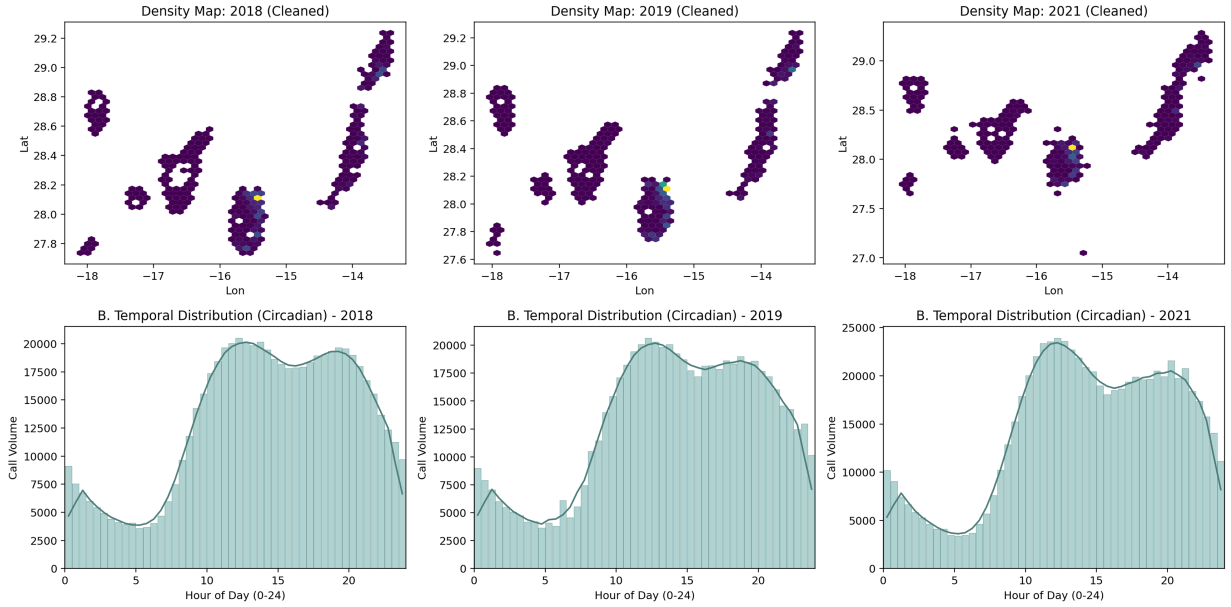

Figure 1: Density map 2018–2019–2021 datasets.

For the geographical density maps, an additional display-domain filter was applied to exclude extreme coordinates outside the Canary Islands visual range. This filter affects only the descriptive figures and the corresponding figure-support count reported in Table 3; it was not used as a duplicate-removal or missing-position cleaning step in the experimental evaluation.

2.2.3 Synthetic TDoS Dataset and TDoS Attack

Given the unavailability of public datasets containing real-world TDoS attacks, a Python-based pipeline was developed to generate and optimize a synthetic dataset consistent with a TDoS scenario targeting telephony infrastructure. The generation process follows the general principles of synthetic denial-of-service traffic construction described in [11–13] while adapting them to the operational characteristics of Public Safety Answering Points (PSAPs).

TDoS attacks aim to exhaust the available human and telecommunication resources of a call center through a large volume of coordinated calls. Unlike volumetric DDoS attacks against IP infrastructure, TDoS attacks specifically target service availability at the telephony and operator level. In the context addressed in this study, the attack model assumes a flood of calls directed at a PSAP, with the objective of keeping operators continuously occupied and delaying or preventing the handling of genuine emergency calls.

The attack model is characterized by temporal and geographical patterns. In the temporal dimension, the relevant properties are burstiness, concentration of calls in short time windows, short call durations compatible with automated behavior, and reduced inter-arrival times between consecutive calls. In the geographical dimension, the model considers both spatial dispersion, which reflects the use of multiple simulated origins, and local accumulations, which emulate coordinated call clusters affecting specific areas.

The reference synthetic TDoS dataset was constructed in two stages. First, a baseline synthetic dataset was generated using the geographical boundaries of the PSAP operational region and applying the following criteria:

- The number of records was defined according to the reference size of the selected baseline data.
- Coordinates were sampled within the polygons of the Canary Islands reference shapefile, preserving geographical validity.
- Spatial concentrations were then introduced by selecting random centers inside the operational area and generating calls around them while maintaining polygon membership
- Dates were distributed throughout the reference year, temporal spikes were inserted at selected minutes, call durations were assigned within operationally plausible ranges, and operators were assigned in a balanced manner to avoid introducing artificial operator bias at this initial stage.

Second, an optimization stage was applied to reinforce TDoS-like properties in a scalable and reproducible way.

- Coordinates were normalized to a common spatial reference when required, dates were converted to a standard *datetime* format, invalid records were removed, and the auxiliary variable *hour* was extracted.
- Temporal burstiness was intensified by increasing the density of calls in the 10:00–12:00 window, preserving chronological order after replication.
- Spatial clustering was then estimated using a scalable procedure based on controlled subsampling, DBSCAN clustering, and nearest-neighbor propagation to the complete dataset. This process generated the binary variable *IS_CLUSTER*, identifying records associated with spatial accumulations.
- Call durations were subsequently adjusted using different distributions for clustered and non-clustered records, and the inter-arrival time variable *delta_t* was compressed for clustered calls to emulate coordinated automated behavior.

The main operations applied during the optimization stage are illustrated in Listings 1 and 2. Listings 1 shows the temporal intensification mechanism used in the original synthetic TDoS dataset. Records falling within the attack window are replicated to create a sustained burst pattern while preserving chronological order.

```
attack_hours = [10, 11, 12]

attack_df = df[df["hour"].isin(attack_hours)]
attack_df = pd.concat([attack_df] * burst_replication_factor,
                      ignore_index=True)

df = pd.concat([df, attack_df], ignore_index=True)
df = df.sort_values("DATE").reset_index(drop=True)
```

Listing 1: Temporal burstiness injection in the original synthetic TDoS dataset

Spatial structure was reinforced through a scalable clustering procedure. Since the complete dataset contains hundreds of thousands of records, clustering is first estimated on a controlled subsample and then propagated to the full dataset using nearest-neighbor search. This produces the binary variable *IS_CLUSTER*, which identifies spatial accumulations compatible with coordinated attack behavior.

```
if len(df) > MAX_CLUSTER_SAMPLE:
    sample_df = df.sample(n=MAX_CLUSTER_SAMPLE,
                          random_state=RANDOM_SEED)
else:
```

```
    sample_df = df.copy()

coords_sample = sample_df[["LATITUDE", "LONGITUDE"]].values

db = DBSCAN(eps=EPS_CLUSTER,
            min_samples=MIN_SAMPLES,
            n_jobs=-1)

labels = db.fit_predict(coords_sample)
sample_df["IS_CLUSTER"] = (labels != -1).astype(int)
df["IS_CLUSTER"] = 0

cluster_points = sample_df.loc[
    sample_df["IS_CLUSTER"] == 1,
    ["LATITUDE", "LONGITUDE"]
].values

if len(cluster_points) > 0:
    nn = NearestNeighbors(n_neighbors=1,
                          algorithm="ball_tree").fit(cluster_points)

    distances, _ = nn.kneighbors(
        df[["LATITUDE", "LONGITUDE"]].values
    )

    df["IS_CLUSTER"] = (distances[:, 0] < EPS_CLUSTER).astype(int)
```

Listing 2: Scalable spatial clustering in the optimized synthetic TDoS dataset

The final synthetic dataset was validated using three criteria: the burstiness index, defined as the peak-to-mean ratio of hourly call counts; the percentage of records flagged as *IS_CLUSTER*; and the geographical consistency of the resulting coordinates. The latter is particularly relevant in this study, since synthetic anomalies must not be trivially detectable as out-of-domain geographical artifacts. Therefore, all synthetic datasets used in the experiments were constrained to the Canary Islands operational domain and formally validated against the reference shapefile. Coordinates falling outside the valid region were repaired by replacing them with valid coordinates sampled from the corresponding scenario whenever possible, or from the validated empirical synthetic support otherwise.

To address the risk of overfitting to a single synthetic-generation logic, the original TDoS dataset was retained as the reference attack scenario and complemented with additional synthetic and benign variants. These variants preserve a size equivalent to the original synthetic dataset and were designed to test the detector under different operational and adversarial conditions. Table 4 summarizes the scenario suite. In this table, Benign high-load scenarios represent non-malicious operational stress conditions. These are not TDoS attacks, but they can legitimately trigger anomaly alerts because their traffic profile deviates from the typical behavior of PSAPs.

Table 4: Synthetic and benign scenarios used in the validation protocol.

Scenario	Scenario Type	Generation Rationale	Validation Purpose
Original synthetic TDoS	Anomaly	Reference attack dataset generated through temporal burst intensification, spatial clustering, short-duration calls, and compressed inter-arrival times.	Reference synthetic attack used as the baseline adversarial scenario

(Continued)

Table 4 (continued)

Scenario	Scenario Type	Generation Rationale	Validation Purpose
Concentrated TDoS	Anomaly	Attack traffic is concentrated in a reduced spatiotemporal region, increasing local density and temporal coordination.	Tests whether the detector identifies highly localized saturation patterns.
Distributed TDoS	Anomaly	Attack traffic is distributed across several valid geographical areas within the Canary Islands support, preserving coordinated temporal behavior.	Evaluates robustness against botnet-like dispersion without relying on a single spatial cluster.
Low-intensity TDoS	Anomaly	Only a fraction of the traffic is modified to include coordinated attack-like behavior, while the remaining records preserve more regular patterns.	Assesses the probability of missed attacks when the malicious signal is less pronounced.
Realistic flash crowd	Benign high-load	A benign emergency-driven surge is generated with high call concentration in time and space, but without malicious automation assumptions.	Evaluates whether the detector confuses legitimate emergency peaks with TDoS attacks.
Non-malicious congestion	Benign high-load	A benign congestion scenario is generated by increasing traffic load and service pressure while preserving legitimate spatial variability.	Tests operational specificity under high-load but non-adversarial conditions.
Zero-day adaptive jitter	Anomaly	Temporal and spatial perturbations are introduced using adaptive jitter, avoiding the fixed burst structure of the original generator.	Tests performance against attack patterns not explicitly encoded in the base generator.
Zero-day replay mimicry	Anomaly	Plausible traces are recombined and shifted in time and space to approximate the empirical support while disrupting normal correlations.	Evaluates vulnerability to mimicry strategies that attempt to resemble legitimate PSAP traffic.

The validation protocol also includes additional synthetic variants designed to avoid dependence on the fixed logic of the original generator. Listing 3 provides a representative example of the zero-day adaptive jitter scenario. In this variant, temporal and spatial perturbations are applied to attack-like records so that the resulting pattern no longer follows the fixed 10:00–12:00 burst structure or a rigid spatial cluster. The purpose is to evaluate whether the detector remains effective when the attack deviates from the assumptions of the original synthetic generator.

```
variant = base_df.copy()

rng = np.random.default_rng(RANDOM_SEED)

attack_mask = rng.random(len(variant)) < zeroday_attack_fraction

# Temporal perturbation: avoid the fixed attack window.
minute_shift = rng.integers(
    low=-zeroday_max_shift_minutes,
    high=zeroday_max_shift_minutes + 1,
    size=attack_mask.sum()
)

variant.loc[attack_mask, "DATE"] = (
    pd.to_datetime(variant.loc[attack_mask, "DATE"]) +
    pd.to_timedelta(minute_shift, unit="m")
)

# Spatial perturbation: introduce adaptive local jitter.
variant.loc[attack_mask, "LATITUDE"] += rng.normal(
    loc=0.0,
    scale=zeroday_lat_jitter,
    size=attack_mask.sum()
)

variant.loc[attack_mask, "LONGITUDE"] += rng.normal(
    loc=0.0,
    scale=zeroday_lon_jitter,
    size=attack_mask.sum()
)

variant = variant.sort_values("DATE").reset_index(drop=True)
```

Listing 3: Representative generation of a zero-day-like adaptive jitter variant

All generated variants were subsequently validated against the Canary Islands reference shapefile. This step ensures that the detector is not evaluated on synthetic records that fall outside the operational geographical domain. When a generated coordinate falls outside the valid region, it is replaced by a valid coordinate sampled from the same scenario whenever possible, or from the validated empirical synthetic support otherwise.

```
inside_mask = points_inside_canary_shapefile(
    longitude=variant["LONGITUDE"].to_numpy(),
    latitude=variant["LATITUDE"].to_numpy(),
    shapefile=CANARY_SHAPEFILE
)

outside_mask = ~inside_mask

if outside_mask.any():
    valid_support = variant.loc[
        inside_mask,
        ["LATITUDE", "LONGITUDE"]
    ]

    if len(valid_support) == 0:
        valid_support = base_df.loc[
            base_inside_mask,
            ["LATITUDE", "LONGITUDE"]
        ]

```

```
    replacement = valid_support.sample(
        n=outside_mask.sum(),
        replace=True,
        random_state=RANDOM_SEED
    ).reset_index(drop=True)

    variant.loc[outside_mask, "LATITUDE"] = replacement["LATITUDE"].values
    variant.loc[outside_mask, "LONGITUDE"] = replacement["LONGITUDE"].values
```

Listing 4: Geographical validation and repair of synthetic coordinates

Listings 3 and 4 summarize the zero-day variant generation and geographical validation procedures. First, zero-day-like variants are generated by disrupting the fixed temporal and spatial assumptions of the original attack generator. Second, every generated scenario is constrained to the Canary Islands operational domain through shapefile-based validation and repair. This design reduces the possibility that the detector learns artifacts of a single generator or classifies records as anomalous simply because they fall outside the valid geographical support.

The main parameters controlling the original optimized synthetic dataset are summarized in Table 5. These parameters regulate reproducibility, temporal intensity, spatial compactness, clustering scalability, call-duration modeling, and temporal coordination.

This extended scenario design allows the model to be evaluated not only against the original synthetic TDoS logic, but also against alternative attack patterns, benign high-load conditions, and zero-day-like variants. As a result, the validation protocol reduces the risk that detection performance is merely the consequence of overfitting to a single synthetic generator.

Table 5: Parameters of the optimized synthetic dataset.

Parameter	Description and Effect
RANDOM_SEED	Seed used to ensure reproducibility in random sampling, reassignment, and scenario generation.
MAX_CLUSTER_SAMPLE	Maximum subsample size used for scalable clustering in large datasets.
EPS_CLUSTER	Neighborhood radius used to identify and propagate spatial clusters.
MIN_SAMPLES	Minimum density required to define a spatial cluster.
attack_hours	Time window in which traffic is intensified to induce burst-like behavior.
burst_replication_factor	Replication factor applied to records within the attack window to increase temporal density.
duration_scale_cluster/ duration_scale_normal	Scale parameters used to model call duration for clustered attack records and regular records.
clip_cluster_range/ clip_normal_range	Operational clipping ranges used to avoid implausible call durations.
top_operators/operator_bias_probs	Operator-selection parameters used to simulate saturation pressure while controlling artificial bias.
delta_t_compression_range	Compression interval applied to *delta_t* in clustered records to reinforce temporal coordination.

2.3 Quantum Machine Learning Concepts

Quantum Machine Learning (QML) integrates the principles of quantum mechanics—such as superposition and entanglement—with classical machine learning to process information in high-dimensional spaces where classical algorithms may struggle [14,15]. In quantum computing, information is represented by qubits, whose state vector in a two-dimensional Hilbert space is denoted as:

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle, \quad \text{where} \quad |\alpha|^2 + |\beta|^2 = 1 \tag{1}$$

Instead of relying on variational approaches [16], the anomaly detection framework proposed in this study utilizes Quantum Kernel methods [17–19]. This requires embedding classical data into a quantum state via a feature map. Recent theoretical work has also shown that several parametrized quantum learning models can be understood within a broader linear-model framework in quantum Hilbert spaces, clarifying the relationship between explicit quantum models, data re-uploading circuits, and kernel-based formulations [20]. This perspective is relevant here because the proposed detector relies on similarities induced by a quantum feature map rather than on an optimized variational ansatz.

To accurately model the interdependencies among the call features (longitude, latitude, and temporal phase), a *PauliFeatureMap* is implemented. Unlike basic encoding methods restricted to individual rotations, this approach utilizes both individual Z-rotations and full-entanglement ZZ-interactions. For a feature set $\{x_1, x_2, x_3\}$ encoded across $n = 3$ qubits, the individual rotation $R_Z(x_i)$ is defined as:

$$R_Z(x_i) = \begin{pmatrix} e^{-ix_i/2} & 0 \\ 0 & e^{ix_i/2} \end{pmatrix} \tag{2}$$

The generated circuit, illustrated in Fig. 2, systematically maps the classical TDoS parameters into the quantum feature space.

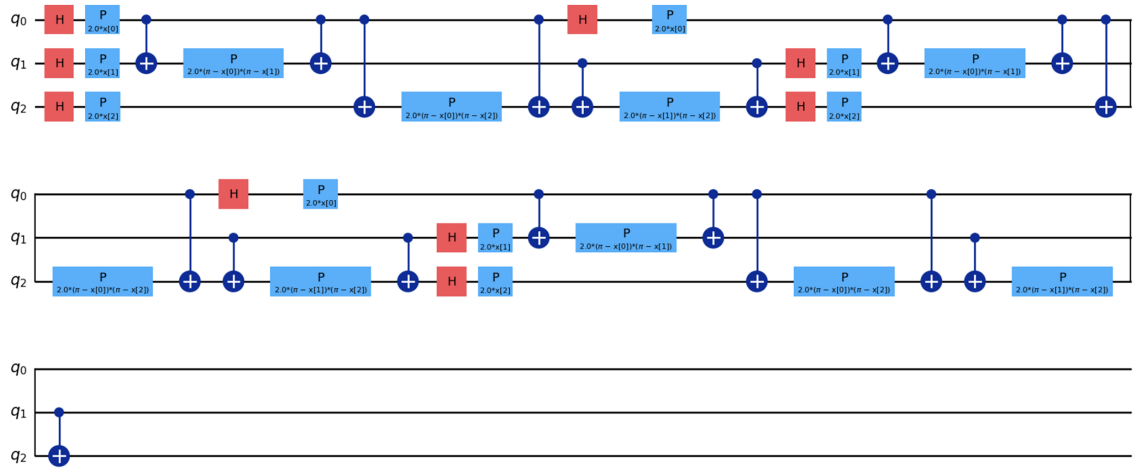

Figure 2: *PauliFeatureMap* circuit applied to 3 qubits featuring Hadamard, Z, and ZZ gates.

The encoding process is executed through a sequence of unitary transformations [21]: First, Hadamard (H) gates are applied to initialize the qubits in an equal superposition of the computational basis states. Second, data-dependent phase rotations (Z gates) encode each classical feature x_i into the phase of the corresponding qubit. Third, to capture high-order non-linear correlations—such as the interaction between a call's geographic origin and its time of occurrence—ZZ gates entangle pairs of qubits based on the product

of their features [22]. Finally, this sequence is repeated across multiple layers (*reps = 3* in this model) to increase the complexity and non-linearity of the quantum feature space.

Following this embedding, the classical call data are mapped into quantum states in an 8-dimensional Hilbert space induced by the 3-qubit PauliFeatureMap. This representation enables the Quantum Kernel k-Nearest Neighbors (QKNN) algorithm [23] to compute similarity metrics and identify volumetric anomalies consistent with potential TDoS attacks.

3 Proposed Framework

3.1 Methodology

This framework defines an unsupervised density-estimation procedure for distinguishing ordinary PSAP call-flow behavior from traffic patterns compatible with Telephony Denial of Service (TDoS) attacks. Classical anomaly-detection models are used as baselines, but direct separation in the original feature space is limited by the structure of PSAP traffic. Legitimate calls may include spatial outliers associated with remote locations or positioning uncertainty, while temporal behavior is affected by circadian rhythms, operational congestion, and emergency-related call surges. In addition, the absence of publicly available real-world TDoS traces requires validation beyond a single synthetic attack generator, including alternative synthetic scenarios and zero-day-like variants.

Given these non-linear characteristics, kernel-based methods provide a suitable mechanism for comparing records through their joint spatiotemporal structure. For an input vector $x \in \mathbb{R}^d$, a feature map $\phi(x)$ embeds the data into a Hilbert space $\boldsymbol{\mathcal{H}}$, where similarities are evaluated through an inner product:

$$K(x, y) = \langle \phi(x), \phi(y) \rangle_{\mathcal{H}}. \tag{3}$$

This formulation makes it possible to evaluate similarity without explicitly constructing the induced feature space. Building upon classical kernel principles [24,25], this study implements a quantum feature map in which a classical data vector x is encoded into a quantum state $|\psi(x)\rangle$ through a parameterized circuit $U(x)$:

$$|\psi(x)\rangle = U(x)|0\rangle^{\otimes n}. \tag{4}$$

The similarity between two records is then measured as the fidelity between their corresponding quantum states:

$$K_Q(x, y) = |\langle \psi(x) | \psi(y) \rangle|^2. \tag{5}$$

This similarity metric supports the local density-estimation stage. The model does not assign malicious intent to individual calls. Instead, it evaluates whether a set of calls deviates from the learned spatiotemporal normality profile. This distinction is important in PSAP environments, where benign emergency peaks may also produce statistically anomalous traffic.

To reduce dependence on a single calibration year, the evaluation extends the original 2019 baseline to include single-year baselines from 2018 and 2021, together with a multiyear baseline combining 2018, 2019, and 2021. Validation is also performed across multiple geographically constrained scenarios: concentrated TDoS, distributed TDoS, low-intensity TDoS, benign flash crowd, non-malicious congestion, and two zero-day-like variants based on adaptive jitter and replay mimicry. In addition, the proposed detector is compared with classical anomaly-detection methods, including Isolation Forest, One-Class SVM, Local Outlier Factor, Gaussian Mixture Models, and classical k-nearest neighbors, under the same sampling conditions.

3.2 Implementation of the TDoS Detection Model

Each call is represented by a three-dimensional vector composed of longitude, latitude, and temporal phase. The temporal feature is transformed into a circular variable normalized in $[0, 2\pi]$, preserving daily periodicity and avoiding artificial discontinuities around midnight.

The detection pipeline uses a dual-kernel architecture. The first modality, the Qiskit-Simulated Quantum Kernel, implements a 3-qubit PauliFeatureMap with single-qubit phase rotations and entangling ZZ interactions. Since the input vector contains three variables, the induced Hilbert space has dimension $2^3 = 8$. This modality provides a simulated quantum-kernel reference, although its execution time is limited by the overhead of classical quantum-circuit simulation.

The use of quantum kernels in near-term settings must nevertheless be interpreted with caution. Their practical advantage depends on the structure of the data, the number of measurements, the effect of noise, and the finite-sample regime in which the kernel is estimated [26]. For this reason, the Qiskit-Simulated Quantum Kernel is used in this work as a methodological reference, while the Quantum-Inspired Kernel provides a computationally practical counterpart for large-scale screening.

The second modality, the Quantum-Inspired Kernel, uses a classical harmonic expansion of the same input features. It combines sine and cosine components at different frequencies to preserve a kernel-based similarity structure while avoiding quantum-circuit simulation. Under the tested conditions, this implementation provides the more computationally practical option for large-scale screening.

Both modalities use the same local density-estimation logic. A reference dictionary of normal calls is sampled from the selected baseline, the affinity of test samples against this reference set is computed, and a normality score is derived from local kernel density. Dataset-level anomalies are identified by comparing the average test density against the calibrated baseline distribution through a Z-Score. The main experiments use $Z > 2.0$ as the operational alert threshold, complemented by a threshold-sensitivity analysis to quantify the trade-off between false positives and false negatives. This threshold was selected through exploratory empirical calibration over the available experimental corpus. It should therefore be interpreted as the default operating point used in this study, rather than as a theoretical cutoff derived from the kernel formulation or as a deployment-independent constant.

Processing times are recorded for both kernel modalities. This comparison provides a practical assessment of computational feasibility: the Qiskit-simulated implementation serves as a methodological reference for the quantum-kernel formulation, whereas the quantum-inspired implementation represents the near-term operational alternative under current computational constraints.

3.2.1 Dual Kernel Architecture

Mapping the three-dimensional call representation into a non-linear similarity space is achieved through two complementary kernel modalities implemented within the same detection pipeline:

- *Qiskit-Simulated Quantum Kernel*. This modality constructs a *PauliFeatureMap* parameterized by the input features: longitude, latitude, and temporal phase. Since the input vector contains three variables, the circuit uses three qubits and induces an 8-dimensional Hilbert space ($2^3 = 8$). Single-qubit phase rotations and entangling ZZ interactions are used to encode spatial–temporal interactions among the call features. The similarity between two input vectors, $\vec{x}$ and $\vec{y}$, is evaluated as the fidelity between their corresponding quantum states:

$$K_Q(\vec{x}, \vec{y}) = |\langle \psi(\vec{x}) | \psi(\vec{y}) \rangle|^2. \tag{6}$$

This modality provides a simulated quantum-kernel reference for the proposed formulation. Since the circuit is evaluated through classical simulation, it is also used to quantify the computational overhead associated with Qiskit-based quantum-kernel evaluation.

- *Quantum-Inspired Kernel*. To preserve a non-linear similarity structure at lower computational cost, a classical harmonic kernel is introduced. The original vector is mapped into sine and cosine components at fundamental and second-order frequencies:

$$\phi(\vec{x}) = [\cos(\vec{x}), \sin(\vec{x}), \cos(2\vec{x}), \sin(2\vec{x})]. \tag{7}$$

Call similarity is then obtained from the normalized inner product of these harmonic expansions. A power transformation, governed by the hyperparameter P (*KERNEL_POWER*), sharpens the distinction between high- and low-affinity records:

$$K_{\mathrm{QI}}(\vec{x}, \vec{y}) = \left(\frac{\phi(\vec{x}) \cdot \phi(\vec{y})^T}{2d} \right)^{2P}, \tag{8}$$

where $d = 3$ is the number of input features. The resulting Gram matrix summarizes the spatiotemporal affinity structure of the dataset, with values close to one indicating strong similarity.

To prevent numerical instability in the logarithmic density estimates and reduce the effect of isolated legitimate outliers in the PSAP data, both kernels are stabilized with a regularization factor λ (*LAMBDA_OPT*):

$$K' = (1 - \lambda)K + \lambda. \tag{9}$$

3.2.2 Local Density Estimation

The computed similarities support an unsupervised local density estimation process. Rather than classifying individual calls as malicious or benign, the method evaluates whether an aggregated set of calls departs from the learned normality profile.

Following a kernel-based nearest-neighbor formulation [23,27], the model identifies the k reference samples from the baseline dictionary with the highest affinity to a given sample $\vec{x}$. The normality score is defined as the logarithmic average of these top affinities:

$$S(\vec{x}) = \log\left(\frac{1}{k} \sum_{j=1}^{k} K(\vec{x}, \vec{x}_j) + \epsilon \right), \tag{10}$$

where K denotes either the Qiskit-simulated quantum kernel or the Quantum-Inspired Kernel, and ϵ is included for numerical stability. High local affinity indicates strong support from the baseline distribution, whereas low affinity indicates weaker spatiotemporal coherence. Sample-level scores are then aggregated into a mean density score for each evaluated dataset and compared with the calibrated baseline distribution.

3.2.3 Normality Characterization and Anomaly Detection

Historical PSAP traffic is used to calibrate the normality profile. The experiments use 2019 as the primary baseline and also incorporate 2018, 2021, and a multiyear baseline combining 2018, 2019, and 2021. This design allows robustness to be assessed across different normal traffic profiles and avoids relying on a single calibration year.

For each configuration, resampling is used to estimate the mean μ_{stable} and standard deviation σ_{stable} of the expected normality scores. New datasets are evaluated by comparing their average density score μ_{test} against this calibrated distribution through a Z-Score:

$$Z = \frac{\mu_{\text{stable}} - \mu_{\text{test}}}{\sigma_{\text{stable}}}. \tag{11}$$

The dataset-level decision rule is then defined as:

$$\hat{y}_d = \begin{cases} 1, & Z_d > \tau, \\ 0, & Z_d \leq \tau, \end{cases}$$

where Z_d denotes the standardized deviation score of dataset d, and τ is the alert threshold.

A positive and sufficiently high Z indicates lower spatiotemporal coherence with respect to the historical baseline. This behavior is expected in coordinated TDoS-like traffic, where temporal and spatial call structures differ from ordinary operations, even when all coordinates remain constrained to the Canary Islands shapefile.

The primary experiments use $\tau = 2.0$ as the default operational alert threshold. This value was not derived analytically from the quantum kernel or from an assumed parametric distribution of the scores. Instead, it was selected through exploratory empirical calibration, in which several candidate Z-score cutoffs were examined to assess the balance between sensitivity to TDoS-like deviations and robustness against false alarms.

Accordingly, $Z > 2.0$ should be interpreted as the calibrated operating point adopted for the experiments reported in this study, not as a universal decision boundary. A threshold-sensitivity analysis was therefore included to evaluate how false positives, false negatives, true-positive rates, and precision change when the decision boundary is made more or less conservative.

Anomalies are flagged at dataset level. Historical years are expected to remain below the threshold, whereas synthetic TDoS and zero-day-like variants are expected to exceed it. Benign high-load scenarios are included as operational stress tests: alerts in these cases are interpreted as benign anomaly alerts rather than as evidence of malicious TDoS activity. Cases in which benign flash crowds or non-malicious congestion exceed the threshold are interpreted as statistical anomalies rather than evidence of malicious intent, and must be assessed together with the operational context available to the emergency service.

4 Results

4.1 Experimental Design

The experimental protocol was designed to evaluate whether the proposed detector can distinguish ordinary PSAP call-flow behavior from traffic patterns compatible with Telephony Denial of Service (TDoS) attacks. The evaluation was conducted using the two kernel modalities described above: the *Quantum-Inspired Kernel*, based on a classical harmonic expansion, and the *Qiskit-Simulated Quantum Kernel*, based on a 3-qubit *PauliFeatureMap*.

The experimental protocol uses 2019 as the reference baseline and extends the analysis to additional single-year baselines, namely 2018 and 2021, together with a multiyear baseline combining 2018, 2019, and 2021. This design makes it possible to assess whether the detector depends on a single calibration year or remains stable under alternative normality profiles.

For each baseline configuration, a reference dictionary of normal calls was sampled and used to estimate local kernel density. Resampling was then applied to characterize the mean and standard deviation of the baseline density scores, denoted by μ_{stable} and σ_{stable}. During inference, each evaluated dataset was assigned an average density score μ_{test}, which was compared with the calibrated baseline using the Z-Score defined in Eq. (11). A positive Z-Score indicates that the evaluated dataset has lower local density than the calibrated normal profile and, consequently, weaker spatiotemporal coherence with ordinary PSAP traffic.

The primary experiments used $Z > 2.0$ as the operational alert threshold. This threshold was not treated as a universal constant; instead, a threshold-sensitivity analysis was conducted to quantify the trade-off between false positives and false negatives.

The evaluation considered three groups of datasets. First, historical yearly datasets from 2011 to 2021 were used to assess false positives on ordinary PSAP traffic, excluding the years used for calibration in each protocol. Second, the original synthetic TDoS dataset was evaluated together with additional synthetic attack variants: concentrated TDoS, distributed TDoS, low-intensity TDoS, and two zero-day-like variants based on adaptive jitter and replay mimicry. Third, benign high-load scenarios, including a realistic flash crowd and non-malicious congestion, were included to evaluate whether the detector reacts to strong but non-adversarial deviations from ordinary traffic. All synthetic scenarios were geographically constrained to the Canary Islands operational domain before evaluation.

Finally, the kernel-based detector was compared with classical anomaly-detection methods under the same sampling conditions. The classical baselines included Isolation Forest, One-Class SVM, Local Outlier Factor, Gaussian Mixture Models, and classical k-nearest neighbors.

4.2 Synthetic, Benign, and Zero-Day-Like Scenarios

Table 6 summarizes the results obtained for the synthetic and benign scenario suite. The table aggregates the behavior of both kernel modalities across the evaluated baseline protocols.

Table 6: Detection results for synthetic, benign, and zero-day-like scenarios.

Kernel	Scenario	Reference Class	Label-Consistent	Label-Inconsistent	$\bar{Z}$	$Z_{\min}$	$Z_{\max}$
QS	Flash crowd	Benign high-load	0	4	47.23	31.13	62.77
QS	Non-malicious congestion	Benign high-load	0	4	19.88	12.62	27.25
QS	Concentrated TDoS	Anomaly	4	0	22.11	4.31	41.91
QS	Distributed TDoS	Anomaly	4	0	19.02	9.57	26.58
QS	Low-intensity TDoS	Anomaly	4	0	25.61	21.01	30.40
QS	Original TDoS	Anomaly	4	0	25.16	20.64	31.49
QS	Zero-day adaptive jitter	Anomaly	4	0	25.76	22.75	29.78
QS	Zero-day replay mimicry	Anomaly	4	0	25.03	22.27	27.15

(Continued)

Table 6 (continued)

Kernel	Scenario	Reference Class	Label-Consistent	Label-Inconsistent	$\bar{Z}$	Z_{min}	Z_{max}
QI	Flash crowd	Benign high-load	0	4	78.00	8.50	123.45
QI	Non-malicious congestion	Benign high-load	0	4	21.07	2.71	31.88
QI	Concentrated TDoS	Anomaly	3	1	8.87	0.68	17.19
QI	Distributed TDoS	Anomaly	4	0	23.18	2.44	40.05
QI	Low-intensity TDoS	Anomaly	4	0	29.27	3.25	46.57
QI	Original TDoS	Anomaly	4	0	28.94	4.41	44.69
QI	Zero-day adaptive jitter	Anomaly	4	0	33.19	5.53	50.40
QI	Zero-day replay mimicry	Anomaly	4	0	32.51	5.40	50.81

Note: QS denotes the Qiskit-Simulated Quantum Kernel and QI denotes the Quantum-Inspired Kernel. The reference class identifies the scenario category used to interpret consistency counting. Label-consistent and label-inconsistent denote the number of baseline protocols in which the detector output agrees or disagrees with the expected behavior for that scenario category. Flash crowd and non-malicious congestion are benign high-load scenarios and are therefore non-adversarial reference cases. Their label-inconsistent outcomes correspond to anomaly alerts, which should be interpreted as benign operational anomaly alerts rather than as confirmed TDoS detections. $\bar{Z}$, $Z_{\min}$, and $Z_{\max}$ denote the mean, minimum, and maximum Z-Score values across protocols.

Table 6 summarizes the consistency between the detector output and the reference class assigned to each synthetic, benign high-load, and zero-day-like scenario.

The Qiskit-Simulated Quantum Kernel is label-consistent for all attack-labeled synthetic scenarios across the four baseline protocols, including the reference TDoS dataset, the concentrated and distributed variants, the low-intensity attack, and the two zero-day-like variants. The Quantum-Inspired Kernel obtains comparable results, with one label-inconsistent outcome in the concentrated TDoS scenario.

The benign high-load scenarios require a separate interpretation. Both the realistic flash crowd and the non-malicious congestion scenario are non-adversarial reference cases, but they exceed the anomaly threshold in all protocols for both kernel modalities. This explains the four label-inconsistent outcomes reported for each of these scenarios. The result indicates sensitivity to operationally relevant departures from ordinary PSAP traffic, but it does not establish malicious intent. In deployment, such alerts would require correlation with contextual information from the emergency service before any mitigation action is taken.

The zero-day-like scenarios address the risk of overfitting to the reference synthetic generator. Both adaptive jitter and replay-mimicry variants are label-consistent across the two kernel modalities, indicating that the detector is not limited to the fixed temporal and spatial assumptions of the reference TDoS construction.

4.3 Threshold Sensitivity

The threshold-sensitivity analysis was included to make the calibration process explicit. Since the default alert threshold was selected empirically, reporting only one operating point would provide an incomplete view of the detector's behavior. Table 7 therefore reports both the default threshold and a more conservative

operating point for each kernel modality, showing how the balance between early detection and false-alarm control changes as the decision boundary is adjusted.

Table 7: Operating points for the Z-Score threshold sensitivity analysis.

Kernel	Z_0	FPR_0	FNR_0	Z^*	FPR	FNR	TPR	Prec.
Qiskit-Simulated	2.00	23.91%	0.00%	4.30	17.39%	0.00%	100.00%	75.00%
Quantum-Inspired	2.00	21.74%	4.17%	2.80	15.22%	8.33%	91.67%	75.86%

Note: Z_0 denotes the default threshold retained after exploratory empirical calibration and used in the main experiments. Z^* denotes a secondary, more conservative operating point selected to illustrate the effect of reducing false positives under the same experimental corpus. FPR, FNR, TPR, and precision should therefore be interpreted as retrospectively calibrated operating estimates, not as prospective deployment guarantees.

Table 7 confirms that the Z-score threshold controls a genuine operating trade-off. For the Qiskit-Simulated Quantum Kernel, increasing the threshold from $Z = 2.0$ to $Z = 4.3$ reduces the false-positive rate from 23.91% to 17.39% without increasing the observed false-negative rate in the evaluated scenarios. For the Quantum-Inspired Kernel, increasing the threshold from $Z = 2.0$ to $Z = 2.8$ reduces the false-positive rate from 21.74% to 15.22%, but increases the false-negative rate from 4.17% to 8.33%. This behavior illustrates the practical role of threshold calibration: lower thresholds favor early detection, whereas higher thresholds reduce false alarms at the cost of potentially missing weaker or more localized attack patterns. Consequently, $Z > 2.0$ should be understood as the default early-warning configuration used in this study, not as a universal threshold.

4.4 Confusion Matrices and Confidence Intervals

Table 8 reports the confusion-matrix counts for each kernel modality and baseline protocol, together with Wilson 95% confidence intervals for the observed false-positive and false-negative rates.

Table 8: Confusion-matrix summary by kernel and baseline protocol with Wilson 95% confidence intervals.

Kernel	Protocol	TN	FP	FN	TP	FPR [95% CI]	FNR [95% CI]
Qiskit-Simulated	Baseline 2018	10	2	0	6	16.67% [4.70, 44.80]	0.00% [0.00, 39.03]
Qiskit-Simulated	Baseline 2019	10	2	0	6	16.67% [4.70, 44.80]	0.00% [0.00, 39.03]
Qiskit-Simulated	Baseline 2021	7	5	0	6	41.67% [19.33, 68.05]	0.00% [0.00, 39.03]
Qiskit-Simulated	Multiyear 2018–2019–2021	8	2	0	6	20.00% [5.67, 50.98]	0.00% [0.00, 39.03]
Quantum-Inspired	Baseline 2018	9	3	0	6	25.00% [8.89, 53.23]	0.00% [0.00, 39.03]
Quantum-Inspired	Baseline 2019	10	2	0	6	16.67% [4.70, 44.80]	0.00% [0.00, 39.03]
Quantum-Inspired	Baseline 2021	9	3	0	6	25.00% [8.89, 53.23]	0.00% [0.00, 39.03]
Quantum-Inspired	Multiyear 2018–2019–2021	8	2	1	5	20.00% [5.67, 50.98]	16.67% [3.01, 56.35]

Note: TN, FP, FN, and TP denote true negatives, false positives, false negatives, and true positives, respectively. FPR is the false-positive rate and FNR is the false-negative rate. Wilson confidence intervals are reported in percentage points. The multiyear protocol uses 2018, 2019, and 2021 as the calibration baseline. Benign high-load scenarios are counted as negative cases for the purpose of the binary confusion matrix, but alerts on these scenarios are interpreted separately as benign operational anomaly alerts.

Table 8 shows that the Qiskit-Simulated Quantum Kernel detects all attack-labeled scenarios in each evaluated baseline protocol, with no observed false negatives. Its main limitation is the presence of false positives, especially under the 2021 baseline, where the observed FPR rises to 41.67%. The Quantum-Inspired

Kernel shows a comparable pattern across the single-year baselines, although it produces one false negative under the multiyear baseline.

The Wilson intervals are included to avoid overinterpreting point estimates obtained from a limited number of historical and synthetic scenarios. In particular, an observed FNR of 0.00% should not be read as evidence that the true false-negative probability is zero; it only indicates that no false negatives were observed under the corresponding experimental protocol. Similarly, the confidence intervals around the false-positive rates reflect the uncertainty associated with the finite number of normal and benign test cases. This uncertainty is particularly important because the decision threshold was calibrated on the available corpus; consequently, the observed error rates should be read as conditional on both the selected threshold and the evaluated scenario set.

The baseline-specific results also indicate that the calibration year influences the balance between sensitivity and specificity. The 2021 baseline is less stable for the Qiskit-Simulated Quantum Kernel, while the multiyear baseline reduces dependence on a single annual normality profile. This supports the use of multiyear calibration as a more conservative configuration for operational analysis.

Fig. 3 provides a visual comparison of the two kernel modalities under the multiyear baseline. This configuration is particularly relevant because it reduces dependence on a single calibration year and provides a more conservative normality profile.

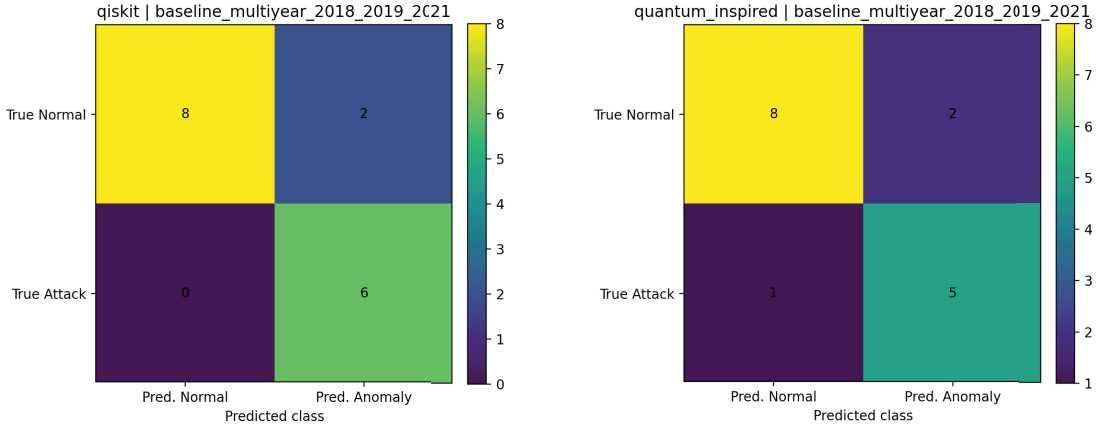

(a) Qiskit-Simulated Quantum Kernel

(b) Quantum-Inspired Kernel

Figure 3: Simplified confusion matrices under the multiyear baseline combining 2018, 2019, and 2021.

Under the multiyear baseline, both kernels produce the same number of false positives among normal or benign cases. The Qiskit-Simulated Quantum Kernel detects all attack-labeled scenarios, whereas the Quantum-Inspired Kernel misses one attack-labeled case. This result illustrates the trade-off between methodological fidelity and computational practicality: the Qiskit-simulated modality provides slightly higher sensitivity in this configuration, while the quantum-inspired modality remains considerably more efficient computationally.

4.5 Comparison with Classical Anomaly-Detection Methods

For empirical comparison with standard baselines, the protocol included several classical anomaly-detection models evaluated under the same sampling conditions. Table 9 reports the results.

Table 9: Comparison with classical anomaly-detection methods by baseline protocol.

Method	Baseline	TP	FP	TN	FN	FPR	FNR	TPR	Prec.	$\bar{Z}$
Classical k-Nearest Neighbors	2018	6	3	9	0	25.00%	0.00%	100.00%	66.67%	15.38
Classical k-Nearest Neighbors	2019	6	2	10	0	16.67%	0.00%	100.00%	75.00%	4.50
Classical k-Nearest Neighbors	2021	6	2	10	0	16.67%	0.00%	100.00%	75.00%	10.09
Classical k-Nearest Neighbors	Multiyear 2018–2019–2021	6	3	7	0	30.00%	0.00%	100.00%	66.67%	9.88
Gaussian Mixture Model	2018	6	3	9	0	25.00%	0.00%	100.00%	66.67%	41.91
Gaussian Mixture Model	2019	6	3	9	0	25.00%	0.00%	100.00%	66.67%	23.54
Gaussian Mixture Model	2021	6	2	10	0	16.67%	0.00%	100.00%	75.00%	27.17
Gaussian Mixture Model	Multiyear 2018–2019–2021	6	2	8	0	20.00%	0.00%	100.00%	75.00%	20.02
Isolation Forest	2018	6	2	10	0	16.67%	0.00%	100.00%	75.00%	10.16
Isolation Forest	2019	6	4	8	0	33.33%	0.00%	100.00%	60.00%	13.84
Isolation Forest	2021	6	4	8	0	33.33%	0.00%	100.00%	60.00%	8.23
Isolation Forest	Multiyear 2018–2019–2021	6	3	7	0	30.00%	0.00%	100.00%	66.67%	12.66
Local Outlier Factor	2018	5	2	10	1	16.67%	16.67%	83.33%	71.43%	9.61
Local Outlier Factor	2019	6	2	10	0	16.67%	0.00%	100.00%	75.00%	15.17
Local Outlier Factor	2021	6	3	9	0	25.00%	0.00%	100.00%	66.67%	6.12
Local Outlier Factor	Multiyear 2018–2019–2021	4	2	8	2	20.00%	33.33%	66.67%	66.67%	3.05
One-Class Support Vector Machine	2018	4	2	10	2	16.67%	33.33%	66.67%	66.67%	0.70
One-Class Support Vector Machine	2019	3	2	10	3	16.67%	50.00%	50.00%	60.00%	0.52
One-Class Support Vector Machine	2021	6	2	10	0	16.67%	0.00%	100.00%	75.00%	3.76
One-Class Support Vector Machine	Multiyear 2018–2019–2021	6	2	8	0	20.00%	0.00%	100.00%	75.00%	3.67
QKNN Qiskit-Simulated Quantum Kernel	2018	6	2	10	0	16.67%	0.00%	100.00%	75.00%	11.01
QKNN Qiskit-Simulated Quantum Kernel	2019	6	2	10	0	16.67%	0.00%	100.00%	75.00%	12.16
QKNN Qiskit-Simulated Quantum Kernel	2021	6	5	7	0	41.67%	0.00%	100.00%	54.55%	12.83
QKNN Qiskit-Simulated Quantum Kernel	Multiyear 2018–2019–2021	6	2	8	0	20.00%	0.00%	100.00%	75.00%	11.15

(Continued)

Table 9 (continued)

Method	Baseline	TP	FP	TN	FN	FPR	FNR	TPR	Prec.	$\bar{Z}$
QKNN Quantum-Inspired Kernel	2018	6	3	9	0	25.00%	0.00%	100.00%	66.67%	22.03
QKNN Quantum-Inspired Kernel	2019	6	2	10	0	16.67%	0.00%	100.00%	75.00%	13.69
QKNN Quantum-Inspired Kernel	2021	6	3	9	0	25.00%	0.00%	100.00%	66.67%	18.25
QKNN Quantum-Inspired Kernel	Multiyear 2018–2019–2021	5	2	8	1	20.00%	16.67%	83.33%	71.43%	1.86

Note: The multiyear baseline combines 2018, 2019, and 2021. TP, FP, TN, and FN denote true positives, false positives, true negatives, and false negatives. FPR is the false-positive rate, FNR is the false-negative rate, TPR is the true-positive rate, and Prec. is precision. $\bar{Z}$ denotes the mean Z-Score across the evaluated protocol-specific cases.

Table 9 compares the proposed QKNN variants with classical anomaly-detection methods under the same baseline protocols and sampling conditions. The results show that several classical methods are competitive, but their behavior varies across calibration years. Classical k-nearest neighbors and Gaussian Mixture Models detect all attack-labeled scenarios in every baseline protocol, with no observed false negatives, although they also generate false positives. Isolation Forest also detects all attack-labeled scenarios, but its false-positive rate increases under the 2019, 2021, and multiyear baselines.

Local Outlier Factor and One-Class SVM are more sensitive to the selected baseline. LOF produces false negatives under the 2018 and multiyear protocols, while One-Class SVM misses attack-labeled scenarios under the 2018 and 2019 baselines. The Qiskit-Simulated QKNN detects all attack-labeled scenarios in every protocol, although the 2021 baseline produces a higher false-positive rate. The Quantum-Inspired QKNN remains competitive across the single-year baselines, but produces one false negative under the multiyear configuration.

These results indicate that the proposed QKNN variants should not be presented as universally superior to classical anomaly detectors. Their contribution is the formulation of a kernel-based spatiotemporal representation that performs comparably to established anomaly detectors while supporting a direct comparison between simulated quantum and quantum-inspired implementations. The baseline-specific comparison also confirms that calibration-year choice affects both classical and kernel-based detectors, supporting the use of multiyear validation in the proposed protocol.

4.6 Geographical Validation of Synthetic Scenarios

The scenario-generation process was audited to ensure that synthetic records remained within the intended operational domain. This geographical control is methodologically important because if synthetic coordinates fall outside the Canary Islands domain, the detector could classify them as anomalous for trivial spatial reasons rather than because of TDoS-like spatiotemporal structure. The generation process therefore constrains and validates the synthetic scenarios against the Canary Islands shapefile before evaluation. This strengthens the interpretation of the synthetic and zero-day-like results.

4.7 Computational Cost and Scalability

For reproducibility, Table 10 reports the hardware platform and the main execution settings used to obtain the timing results. The Qiskit-Simulated Quantum Kernel was evaluated locally through simulated fidelity computation; no physical quantum processor was involved.

Table 10: Execution environment and effective settings for the computational-cost experiments.

Item	Setting Used in the Experiments
Hardware platform	Local workstation with a Silicon M5 processor and 32 GB RAM. No external GPU or physical quantum processor was used.
Qiskit evaluator/backend	The Qiskit-Simulated Kernel was implemented with `qiskit_machine_learning.kernels.FidelityQuantumKernel` and a `PauliFeatureMap`. No IBM Quantum backend, Aer backend, transpilation target, or shot-based sampler was explicitly instantiated. The reported Qiskit timings therefore correspond to local simulated fidelity evaluation.
Quantum feature map	Three-qubit `PauliFeatureMap` with `feature_dimension = 3`, `reps = 3`, `paulis = ["Z", "ZZ"]`, and `entanglement = "full"`. The induced Hilbert-space dimension is $2^3 = 8$.
Effective sample sizes	Each evaluated dataset uses $N_{\text{ref}} = 200$ baseline reference samples and up to $N_{\text{test}} = 150$ test samples. Thus, a single inference run evaluates up to $N_{\text{test}} \times N_{\text{ref}} = 30{,}000$ pairwise affinities. For the Qiskit-Simulated Kernel, test samples are processed in batches of 25.
Timing scope	Per-dataset time is measured as `preprocess_seconds + inference_seconds`. Baseline calibration, table generation, plotting, and full protocol-level elapsed time are excluded.

Under these conditions, Table 11 reports the processing times obtained for the Quantum-Inspired Kernel and the Qiskit-Simulated Quantum Kernel across historical and synthetic datasets.

Table 11: Processing-time comparison by baseline protocol and dataset group.

Baseline Protocol	Dataset Group	N	QI Time (s)	QI Range (s)	QS Time (s)	QS Range (s)	Overhead
Baseline 2018	Historical years	10	2.78	2.11–3.22	35.92	35.26–36.61	13.12×
Baseline 2018	Synthetic suite	8	0.52	0.41–0.63	36.20	36.09–36.42	70.71×
Baseline 2019	Historical years	10	1.21	0.92–1.42	36.50	36.20–37.32	30.76×
Baseline 2019	Synthetic suite	8	0.52	0.41–0.63	35.27	35.14–35.55	69.75×
Baseline 2021	Historical years	10	1.19	0.91–1.41	35.27	35.08–35.41	30.19×
Baseline 2021	Synthetic suite	8	0.52	0.41–0.63	35.29	35.17–35.38	69.42×

(Continued)

Table 11 (continued)

Baseline Protocol	Dataset Group	N	QI Time (s)	QI Range (s)	QS Time (s)	QS Range (s)	Overhead
Multiyear 2018–2019–2021	Historical years	8	1.21	0.91–1.41	35.32	35.17–35.51	29.67×
Multiyear 2018–2019–2021	Synthetic suite	8	0.52	0.40–0.63	35.33	35.17–35.52	69.65×

Note: QI denotes the Quantum-Inspired Kernel and QS denotes the Qiskit-Simulated Quantum Kernel. QI time and QS time report the mean per-dataset wall-clock time within each dataset group, computed as `preprocess_seconds + inference_seconds`. Preprocessing includes data loading or cache retrieval, coordinate parsing, feature construction, and scaling. Inference corresponds to the QKNN `score_samples` call, including kernel evaluation and local-density scoring. The range columns report the minimum and maximum observed per-dataset values within each group. Overhead is computed as the mean QS time divided by the mean QI time. Baseline calibration, table generation, plotting, and full protocol-level elapsed time are not included in these per-dataset measurements.

The Qiskit-Simulated Quantum Kernel required approximately 35–36 s per evaluated dataset across all baseline protocols. The Quantum-Inspired Kernel was considerably faster, processing the synthetic suite in approximately 0.52 s on average and the historical yearly datasets in approximately 1.19–2.78 s, depending on the baseline protocol.

For the synthetic scenarios, local Qiskit simulation was approximately 69–71 times slower than the Quantum-Inspired Kernel. For historical yearly datasets, the overhead ranged from approximately 13 times under the 2018 baseline to approximately 30 times under the 2019, 2021, and multiyear configurations. These results support the use of the Qiskit-Simulated Quantum Kernel as a methodological reference, while identifying the Quantum-Inspired Kernel as the more practical option for near-term operational screening.

The Gram-matrix construction remains the principal scalability constraint because exact pairwise evaluation scales quadratically with the number of samples. Table 12 reports the theoretical memory requirements and observed kernel computation times for the Quantum-Inspired Kernel.

Table 12: Theoretical memory and observed computation time for Gram-matrix construction.

N	Gram Shape	Pairs	Memory (MB)	Kernel Time (s)
50	50 × 50	2500	0.019	0.0001
100	100 × 100	10,000	0.076	0.0006
200	200 × 200	40,000	0.305	0.0003
500	500 × 500	250,000	1.907	0.0017
1000	1000 × 1000	1,000,000	7.629	0.0060
2000	2000 × 2000	4,000,000	30.518	0.0246

Note: N denotes the number of samples used to construct the exact Gram matrix. Memory is computed assuming a dense float64 matrix. The reported times correspond to the Quantum-Inspired Kernel benchmark. Exact pairwise Gram-matrix construction scales as $O(N^2)$ in both memory and pairwise evaluations.

The Quantum-Inspired Kernel scales consistently with the expected $O(N^2)$ pairwise structure. In the reported benchmark, the theoretical memory requirement increases from approximately 0.019 MB for $N = 50$ to approximately 30.52 MB for $N = 2000$. Although the observed computation times remain

low in this benchmark, larger operational deployments would require windowed processing, reference dictionaries, approximate nearest-neighbor search, or low-rank Gram approximations to avoid quadratic growth becoming prohibitive.

5 Discussion

The proposed QKNN-based detector was evaluated across multiple calibration profiles, diverse synthetic and benign scenarios, zero-day-like variants, and established classical anomaly-detection methods. This experimental design provides a cautious assessment of kernel-based spatiotemporal anomaly detection for TDoS early warning, avoiding dependence on a single reference year or on a single synthetic attack generator.

Most attack-labeled datasets—including the reference TDoS, concentrated, distributed, low-intensity, and zero-day-like variants—were successfully identified by both kernel modalities. Specifically, the Qiskit-Simulated Quantum Kernel registered no observed false negatives under the tested baseline protocols. While the Quantum-Inspired Kernel exhibited comparable performance, a single false negative occurred for a concentrated TDoS scenario when using the multiyear baseline. This difference suggests that the classical harmonic approximation largely retains the discriminative behavior of its simulated counterpart, although it may experience a slight sensitivity loss in specific borderline cases.

Importantly, such detections reflect spatiotemporal deviation rather than verified malicious intent. Because the framework measures departures from learned normality, benign events such as flash crowds or non-malicious congestion can also trigger anomalies. A legitimate, large-scale emergency may alter traffic density in ways that resemble a saturation attack, reinforcing the system's intended function as a decision-support mechanism for PSAP operators rather than an autonomous blocking tool.

Calibration profiles heavily influence this model behavior. Relying exclusively on a single year, such as 2021, can introduce year-specific traffic biases, which manifested as an elevated false-positive rate in the Qiskit implementation. Merging data from 2018, 2019, and 2021 into a multiyear baseline reduces this dependence, providing a conservative and operationally stable reference, albeit at the cost of marginally reduced sensitivity for the Quantum-Inspired Kernel. Furthermore, statistical uncertainty must be factored into the reported error rates. Although several configurations yielded an observed 0.00% false-negative or false-positive rate, these metrics represent specific outcomes within the evaluated constraints, not absolute guarantees. The associated Wilson confidence intervals remain essential for framing these point estimates accurately given the finite number of test cases.

Adjusting the operational alert threshold dictates the final balance between sensitivity and specificity. While $Z > 2.0$ served as the default benchmark, the sensitivity analysis confirmed its role as a tunable parameter. Raising this threshold in the Qiskit model reduced false alarms without compromising detection rates for the evaluated attacks. Conversely, the Quantum-Inspired Kernel faced a sharper trade-off, where minimizing false positives increased the likelihood of missing weaker disruptions. Consequently, threshold selection must align directly with the specific risk tolerance of the emergency service. A related limitation is that the threshold calibration was exploratory and retrospective. The selected operating point was obtained by examining the available historical, benign, synthetic, and zero-day-like scenarios, rather than by fixing the threshold in advance and validating it prospectively on an independent labeled TDoS corpus. This choice was unavoidable given the absence of public real-world TDoS traces, but it constrains the interpretation of the reported false-positive and false-negative rates. In operational deployment, baseline construction, threshold calibration, and final validation should be separated whenever sufficient operator-validated events become available.

Contextualizing these findings alongside classical anomaly detectors clarifies the framework's relative contribution. Methods such as k-nearest neighbors, Gaussian Mixture Models, and Isolation Forest also detected the attack scenarios but generally produced false positives depending on the baseline configuration. Others, such as Local Outlier Factor and One-Class SVM, proved highly sensitive to the calibration data and occasionally missed attacks. The proposed QKNN variants do not claim universal dominance over these established alternatives; rather, they introduce a competitive, kernel-based representation of spatiotemporal coherence that enables comparison between a practical Quantum-Inspired implementation and a Qiskit-simulated quantum-kernel reference.

To verify that anomaly detection was not an artifact of overfitting to the reference generator's fixed parameters, such as the 10:00–12:00 burst window, zero-day-like attacks were introduced. Both replay-mimicry and adaptive jitter variants triggered anomalies, indicating a response to broader spatiotemporal disruption rather than memorized temporal patterns. Additionally, since all synthetic and benign traces were geographically constrained to the Canary Islands and validated against the reference shapefile, the model could not rely on out-of-domain coordinate errors to artificially flag malicious activity.

Operational deployment considerations highlight a clear divergence between the two kernel modalities. The Qiskit implementation provides a methodological reference for the 3-qubit *PauliFeatureMap* formulation, yet its reliance on classical simulation imposes a 35–36 s overhead per dataset. By contrast, the Quantum-Inspired Kernel processes synthetic suites in roughly 0.52 s and historical datasets in 1.19–2.78 s, positioning it as the viable path for near-term implementation.

Scalability, however, remains a fundamental bottleneck. Pairwise kernel evaluation dictates an $\mathcal{O}(N^2)$ growth in both memory and computation. Benchmarks for the harmonic approximation demonstrated dense float64 memory requirements escalating from 0.019 MB (N = 50) to 30.518 MB (N = 2000), alongside an increase from 2500 to 4,000,000 pairwise comparisons. Transitioning from a controlled experimental environment to a real-time streaming architecture will therefore necessitate low-rank approximations such as Nyström methods, sliding windows, or bounded reference dictionaries. A second limitation concerns the statistical behavior of quantum kernels themselves. Recent work has shown that, under certain conditions, quantum-kernel values may concentrate exponentially as the number of qubits increases, particularly in the presence of highly expressive embeddings, global measurements, entanglement, or noise [28]. Although the present model uses only three qubits, this issue becomes relevant for future extensions involving higher-dimensional feature maps or hardware-based execution.

Recognizing the boundaries of this study is necessary for future development. The absence of publicly available, real-world TDoS traces restricts evaluation to synthetic models, which, despite deliberate diversification, cannot fully emulate adversarially optimized attacks. Moreover, distinguishing between malicious saturation and legitimate flash crowds based purely on metadata is not fully identifiable from metadata alone. Addressing this limitation likely requires integrating contextual evidence, such as SIP signaling features, audio-level data, or operator workload metrics. Finally, the Qiskit framework operates strictly within a simulated environment, serving a methodological purpose rather than reflecting current quantum hardware readiness.

Overall, this kernel-based approach provides a structured and interpretable anomaly signal for PSAP environments. Its practical reliability depends on baseline calibration, threshold selection, operational context, and computational constraints. Therefore, the detector should be interpreted as a decision-support component rather than as a stand-alone mechanism for automatically determining malicious activity.

6 Conclusions and Future Work

To address the critical threat of Telephony Denial of Service (TDoS) attacks against Public Safety Answering Points (PSAPs), this study introduces a kernel-based anomaly detection framework. By representing incoming calls through their longitude, latitude, and temporal phase, the system models traffic patterns via an unsupervised Quantum k-Nearest Neighbors (QKNN) local density estimator. This architecture was tested under two distinct modalities: a formal Qiskit-Simulated Quantum Kernel utilizing a 3-qubit *PauliFeatureMap*, alongside a computationally efficient Quantum-Inspired Kernel driven by classical harmonic expansion.

The experimental methodology avoids dependence on single-year calibration. Normality profiles were rigorously constructed and cross-evaluated using data from 2018, 2019, 2021, and a comprehensive multiyear baseline. Furthermore, the evaluation suite expanded well beyond the initial synthetic TDoS reference to incorporate concentrated, distributed, and low-intensity variants, alongside benign flash crowds, non-malicious congestion, and two zero-day-like attack models. Because all synthetic traces were constrained to the Canary Islands' operational domain and validated against regional shapefiles, the framework reduces the risk of trivial detection based merely on out-of-domain geographic coordinates.

Both kernel approaches successfully generated interpretable spatiotemporal anomaly scores. The Qiskit-Simulated modality achieved a 0.00% observed false-negative rate across the attack-labeled scenarios, and the Quantum-Inspired alternative missed only a single concentrated TDoS event under the multiyear baseline. However, the system also flagged non-malicious congestion and benign flash crowds as anomalous. Rather than a flaw, this behavior aligns with the fundamental design of the detector: it quantifies statistical deviations from historical normality, rather than definitively classifying malicious intent. Consequently, this tool is best deployed as an early-warning layer to support PSAP decision-making, not as an automated call-dropping filter.

The inclusion of multiple baselines proved vital for balancing sensitivity and specificity. Relying solely on the 2021 profile, for instance, noticeably inflated the false-positive rate for the Qiskit-Simulated kernel, highlighting the risk of overfitting to year-specific traffic anomalies. Conversely, the multiyear baseline smoothed out these annual idiosyncrasies, offering a more conservative and robust operational profile. When interpreting these outcomes, the derived Wilson confidence intervals remain essential; an observed 0.00% error rate reflects specific experimental conditions rather than an absolute guarantee of zero operational risk. This interpretation also applies to the Z-score threshold itself: $Z > 2.0$ was retained as an empirically calibrated early-warning threshold for the present experiments, and should be recalibrated before use in a different PSAP environment or under substantially different traffic conditions.

When benchmarked against classical machine learning approaches, the proposed QKNN variants proved highly competitive, though they do not universally outclass traditional methods. Established algorithms like Gaussian Mixture Models and classical k-nearest neighbors successfully identified all attack scenarios, yet struggled with elevated false-positive rates. Other methods (One-Class SVM, Local Outlier Factor) exhibited high sensitivity to the chosen calibration profile, resulting in occasional missed attacks. The primary contribution here is not necessarily outright algorithmic supremacy, but rather the successful formulation of a quantum-compatible representation of PSAP traffic that performs competitively against classical baselines, offering a clear comparative bridge between simulated quantum and quantum-inspired implementations.

From a computational standpoint, a clear operational divide separates the two kernel modalities. While the Qiskit-Simulated kernel provides a methodological reference for the quantum-kernel formulation, its

reliance on classical simulation bottlenecks execution, requiring roughly 35 to 36 s per dataset. The Quantum-Inspired Kernel bypasses this limitation, processing synthetic suites in an average of 0.52 s and historical yearly data in under three seconds. Under current hardware and simulation constraints, this positions the harmonic approximation as the most practical candidate for near-term experimentation.

Future work should focus on the following points:

- *Low-latency and streaming detection*. Moving beyond controlled batch evaluations, the architecture must be adapted for continuous streaming environments. Addressing the quadratic scaling of exact pairwise kernel computations will require integrating sliding windows, bounded reference dictionaries, approximate nearest-neighbor searches, or low-rank Gram-matrix approximations.
- *Contextual threshold calibration*. The default $Z > 2.0$ alert threshold serves merely as an initial operational point. Practical deployments require dynamic calibration tailored to the specific risk tolerance of the PSAP, factoring in expected traffic volatility and the tangible costs of false alarms vs. missed attacks. Future work should therefore evaluate prospective calibration protocols in which thresholds are fixed before testing and subsequently updated only through predefined operational rules.
- *Empirical adversarial validation*. The absence of publicly accessible, real-world TDoS datasets limits the current evaluation to synthetic and zero-day-like analogs. Validating the framework against anonymized incident traces, controlled red-team telephony exercises, or simulated operator saturation drills remains a critical next step.
- *Multimodal data integration*. Distinguishing a malicious saturation attack from a legitimate, high-volume emergency cannot rely exclusively on metadata. Future iterations should fuse spatiotemporal density estimates with SIP signaling characteristics, session establishment latencies, network entropy indicators, operator workload metrics, and, where legally permissible, audio-level features.
- *Empirical quantum hardware testing*. The current Qiskit implementation serves a purely methodological purpose via simulation. As physical quantum processors mature, executing the 3-qubit *PauliFeatureMap* on actual quantum hardware will be necessary to benchmark real-world fidelity, gate noise, and latency against this simulated baseline.
- *Adversarial robustness modeling*. Further investigation is needed to determine the framework's resilience against highly sophisticated threat actors. This involves modeling attackers capable of executing replay attacks or adapting their temporal-spatial distribution to closely mimic the PSAP's learned normality profile, potentially eroding the detector's operational margin.

In summary, the proposed framework provides an interpretable anomaly signal for TDoS-compatible deviations in PSAP traffic. Its operational value will depend on prospective threshold calibration, scalable kernel approximation, and integration with contextual evidence from the emergency service. Translating this theoretical viability into practical reliability will depend on rigorous baseline calibration, informed threshold tuning, and scalable infrastructure. Until quantum processing units reach operational maturity, the Qiskit-Simulated model will remain a methodological reference, while the Quantum-Inspired Kernel offers the immediate computational efficiency necessary for real-world screening.

Acknowledgement: The authors would like to thank the Canary Islands PSAP in Spain for providing the anonymized data used in this study.

Funding Statement: This work has been possible thanks to the PID2022-138933OB-I00 ATQUE and the 2023DIG28 IACTA research projects funded by MCIN/AEI/10.13039/501100011033/FEDER EU and the CajaCanarias la Caixa Foundation. It is also part of the Cybersecurity Chair of the University of La Laguna funded by Binter.

Author Contributions: Carlos Rosa-Remedios: Methodology, Formal Analysis and Writing—Original Draft Preparation. Pino Caballero-Gil: Methodology and Writing—Review and Editing. Jezabel Molina-Gil: Investigation and Review and Editing. All authors reviewed and approved the final version of the manuscript.

Availability of Data and Materials: The working dataset is based on anonymized information recorded by the Canary Islands Emergency Service between 2011 and 2021 in the target geographical area of the Canary Islands, obtained upon request from the management of this service, as this data is not publicly available.

Ethics Approval: Not applicable. This study does not involve human or animal participants, and the data used are completely anonymous and are not affected in any way by national and European data protection regulations.

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

References

1. Cerezo M, Verdon G, Huang HY, Cincio L, Coles PJ. Challenges and opportunities in quantum machine learning. Nature Comput Sci. 2022;2(9):567–76. doi:10.1038/s43588-022-00311-3.
2. Rosa-Remedios C, Caballero-Gil P. Optimizing quantum machine learning for proactive cybersecurity. Optimiz Eng. 2025;26(4):2321–53. doi:10.1007/s11081-024-09934-z.
3. National Institute of Standards and Technology. The NIST Cybersecurity Framework (CSF) 2.0. National Institute of Standards and Technology (NIST); 2024. NIST CSWP 29. [cited 2026 Jun 10]. Available from: https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.29.pdf.
4. Guri M, Mirsky Y, Elovici Y. 9-1-1 ddos: threat, analysis and mitigation. arXiv:1609.02353. 2016.
5. National Emergency Number Association. NENA telephony denial of service (TDoS) information document. National Emergency Number Association (NENA); 2025. NENA-INF-045.1-2025. [cited 2026 Jun 10]. Available from: https://www.nena.org/.
6. Kadri MR, Abdelli A, Othman JB, Mokdad L. Survey and classification of Dos and DDos attack detection and validation approaches for IoT environments. Inter Things. 2024;25(4):101021. doi:10.1016/j.iot.2023.101021.
7. McLachlan GJ, Peel D. Finite mixture models. Hoboken, NJ, USA: John Wiley & Sons; 2000.
8. Aronszajn N. Theory of reproducing kernels. Trans Am Mathem Soc. 1950;68(3):337–404. doi:10.1090/s0002-9947-1950-0051437-7.
9. Schölkopf B, Smola AJ. Learning with kernels: support vector machines, regularization, optimization, and beyond. Cambridge, MA, USA: MIT press; 2002.
10. Rodríguez-Díaz F, Gutiérrez-Avilés D, Troncoso A, Martínez-Álvarez F. A survey of quantum machine learning: foundations, algorithms, frameworks, data and applications. ACM Comput Surv. 2025;58(4):1–35.
11. Arnaboldi L, Morisset C. Generating synthetic data for real world detection of DoS attacks in the IoT. In: Federation of International Conferences on Software Technologies: Applications and Foundations. Cham, Switzerland: Springer; 2018. p. 130–45.
12. Febro A, Xiao H, Spring J. Telephony denial of service defense at data plane (TDoSD@ DP). In: Proceedings of the NOMS 2018–2018 IEEE/IFIP Network Operations and Management Symposium; 2018 Apr 23–27; Taipei, Taiwan. p. 1–6.
13. Kumar V, Sinha D. Synthetic attack data generation model applying generative adversarial network for intrusion detection. Comput Secur. 2023;125(9):103054. doi:10.1016/j.cose.2022.103054.
14. Ciliberto C, Herbster M, Ialongo AD, Pontil M, Rocchetto A, Severini S, et al. Quantum machine learning: a classical perspective. Proc Royal Soc A Mathem Phys Eng Sci. 2018;474(2209):20170551. doi:10.1098/rspa.2017.0551.
15. Nielsen MA, Chuang IL. Quantum computation and quantum information. Cambridge, UK: Cambridge University press; 2010.
16. Cerezo M, Arrasmith A, Babbush R, Benjamin SC, Endo S, Fujii K, et al. Variational quantum algorithms. Nat Rev Phys. 2021;3(9):625–44.
17. Liu Y, Arunachalam S, Temme K. A rigorous and robust quantum speed-up in supervised machine learning. Nat Phys . 2021;17(9):1013–7. doi:10.1038/s41567-021-01287-z.

18. Schuld M. Supervised quantum machine learning models are kernel methods. arXiv:2101.11020. 2021.
19. Thanasilp S, Wang S, Nghiem NA, Coles P, Cerezo M. Subtleties in the trainability of quantum machine learning models. Quant Mach Intell. 2023;5(1):21. doi:10.1007/s42484-023-00103-6.
20. Jerbi S, Fiderer LJ, Poulsen Nautrup H, Kübler JM, Briegel HJ, Dunjko V. Quantum machine learning beyond kernel methods. Nat Commun. 2023;14(1):517.
21. Barenco A, Bennett CH, Cleve R, DiVincenzo DP, Margolus N, Shor P, et al. Elementary gates for quantum computation. Phys Rev A. 1995;52(5):3457. doi:10.1103/physreva.52.3457.
22. Horodecki R, Horodecki P, Horodecki M, Horodecki K. Quantum entanglement. Rev Mod Phy. 2009;81(2):865–942. doi:10.1103/revmodphys.81.865.
23. Gong LH, Ding W, Li Z, Wang YZ, Zhou NR. Quantum k-nearest neighbor classification algorithm via a divide-and-conquer strategy. Adv Quant Technol. 2024;7(6):2300221. doi:10.1002/qute.202300221.
24. Cortes C, Vapnik V. Support-vector networks. Mach Learn. 1995;20(3):273–97. doi:10.1007/bf00994018.
25. Steinwart I, Christmann A. Support vector machines. New York, NY, USA: Springer Science & Business Media; 2008.
26. Wang X, Du Y, Luo Y, Tao D. Towards understanding the power of quantum kernels in the NISQ era. Quantum. 2021;5:531.
27. Wiebe N, Kapoor A, Svore KM. Quantum nearest-neighbor algorithms for machine learning. Quant Inform Comput. 2015;15(3–4):318–58.
28. Thanasilp S, Wang S, Cerezo M, Holmes Z. Exponential concentration in quantum kernel methods. Nat Commun. 2024;15(1):5200. doi:10.21203/rs.3.rs-2296310/v1.