

AMSA: An Adaptive ML-Based Synchronization Arbiter for Resilient Time Distribution in 5G and Beyond Networks

Ivelin R. Penkov

Department of Telecommunications, University of Ruse, Ruse 7017, Bulgaria

ipenkov@uni-ruse.bg

Abstract: Fifth-generation (5G) and emerging 6G mobile networks impose stringent timing requirements ($|TE| \leq \pm 1.5 \mu\text{s}$ per ITU-T G.8271.1) that the static Best Master Clock Algorithm (BMCA) of IEEE 1588 cannot reliably guarantee under real-world conditions involving packet delay variation (PDV), Global Navigation Satellite System (GNSS) outages, and signal spoofing. This paper presents AMSA (Adaptive ML-based Synchronization Arbiter), a four-stage intelligent software overlay that replaces the deterministic BMCA with a data-driven, context-aware time source selection mechanism. The proposed architecture integrates Gamma-distributed PDV filtering, polynomial-regression-based holdover prediction (termed Virtual Rubidium), trust-score fault detection, and a Random Forest classifier trained on a hybrid dataset of operational telemetry and synthetic anomalies for adaptive source arbitration. AMSA is validated through a Monte Carlo simulation framework across four distinct operational scenarios and corroborated with field measurement data from an operational 5G transport network. Results demonstrate that AMSA reduces the 95th-percentile time error by 55% compared to standard BMCA, eliminates all 209 ITU-T G.8271.1 mask violations observed under high-PDV conditions, extends holdover autonomy by over 30%, and detects covert GNSS spoofing within 250 seconds. Field validation across three scenarios achieves 100% agreement with independent expert assessments, including correct classification of out-of-distribution anomalies.

Received 15 November 2024

Accepted 16 December 2024

Published 02 February 2025

Keywords: 5G synchronization, IEEE 1588, BMCA, machine learning, random forest, precision time protocol, time error, holdover, GNSS spoofing, ITU-T G.8271.1.

1. Introduction

The transition from fourth- to fifth-generation mobile networks has fundamentally altered the role of time synchronization in telecommunications infrastructure. While 4G Long Term Evolution (LTE) systems tolerated phase alignment budgets of several microseconds, 5G New Radio (NR) in Time Division Duplex (TDD) mode mandates absolute time error below $\pm 1.5 \mu\text{s}$ across the entire timing distribution chain, from the Primary Reference Time Clock (PRTC) to the base station antenna port, as specified by ITU-T Recommendation G.8271.1 [1]. This requirement arises from the physical-layer sensitivity of TDD frame alignment, where timing errors exceeding the guard interval cause inter-slot interference and cell outages.

Emerging network capabilities impose tighter constraints. Ultra-Reliable Low-Latency Communications (URLLC) services require deterministic latency budgets that presuppose sub-microsecond synchronization [2]. Time-Sensitive Networking (TSN) integration with 5G, as

defined in 3GPP Release 16, demands end-to-end time awareness across heterogeneous domains [3]. Coordinated Multipoint (CoMP) transmission requires inter-site synchronization well below $\pm 1 \mu\text{s}$ for constructive interference [1, 4].

The prevailing mechanism for time source selection is the Best Master Clock Algorithm (BMCA) of IEEE 1588 PTP [4]. BMCA compares declared clock attributes—class, accuracy, variance, and priority fields—selecting the highest-ranking source. While effective in controlled environments, BMCA exhibits three fundamental limitations in operational 5G transport networks. First, the algorithm suffers from blindness to measured quality; it evaluates self-reported attributes rather than actual performance, potentially selecting a 'GPS-locked' source (clockClass 6) that is actually delivering degraded timing due to multipath or PDV accumulation. Second, the absence of decision memory renders each evaluation cycle independent, causing oscillatory switching (flapping) between sources of similar priority, which introduces transient phase errors of 50–200 ns at each transition. Third, the reliance on static holdover means that upon reference loss, the system freezes the last frequency correction without compensating for second-order drift,



© The Author(s). Published by CIEES Association Press. This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 4.0 Generic License.

leading to quadratic error accumulation in TCXO oscillators [5, 6].

Several approaches address isolated aspects: Kalman filters improve noise rejection but assume Gaussian PDV; fuzzy controllers introduce qualitative reasoning but require expert-crafted rulesets; deep reinforcement learning methods are computationally intensive and lack explainability—a critical requirement for telecom operators who must audit network decisions [5, 7, 8].

This paper presents AMSA (Adaptive ML-based Synchronization Arbiter), a four-stage software overlay. The key contributions are: (1) a complete architecture operating on existing ITU-T G.8275.1 infrastructure without hardware modifications; (2) Gamma-based PDV filtering eliminating mask violations under extreme stress; (3) a Virtual Rubidium concept extending holdover by 30%+ using TCXO hardware; and (4) a validation framework grounded in a hybrid training strategy—merging real-world telemetry with synthetic anomalies—corroborated by field measurements.

The remainder is organized as follows: Section II reviews related work; Section III presents the architecture; Section IV details the hybrid data methodology and simulation framework; Section V reports experimental results; Section VI concludes.

2. Related works

Clock synchronization over packet networks has been surveyed by Lévesque and Tipper [5], who identify PDV as the primary impairment and categorize mitigation strategies into hardware-based (boundary clocks, transparent clocks) and software-based (filtering, servo design) approaches. Purely hardware approaches become prohibitively expensive at the scale required by 5G densification [5].

2.1. Kalman Filter Approaches.

Giorgi proposed an event-based Kalman filter that dynamically adjusts measurement noise covariance based on observed PDV statistics. The approach achieves excellent performance when the noise model matches Gaussian assumptions. However, field measurements in loaded IP/MPLS networks consistently reveal non-Gaussian, heavy-tailed delay distributions which violate this fundamental assumption and lead to suboptimal filter tuning and potential servo instability under real traffic conditions [7, 5, 9, 12].

Specifically, under Gamma-distributed PDV with shape parameter $\alpha=0.5$, the heavy tail produces outliers reaching 2–3 μs that corrupt the Kalman state estimate. The filter’s optimal gain, calibrated for the assumed Gaussian variance, systematically underweights these extreme events, causing the servo to track corrupted timing signals as if they were legitimate phase corrections.

2.2. Disturbance Compensation Methods.

Seo et al. introduced ARMA-based disturbance compensation, improving tracking under non-stationary conditions. While effective at the servo level, this approach requires per-deployment parameter tuning and provides no source selection logic—it cannot determine which source to track, only how to track better [8].

2.3. ML-Driven Approaches.

Wang et al. proposed intelligent time synchronization for 5G-TSN integration, addressing multi-gNB competition. Chandramouli et al. outlined the evolution toward 6G timing services, identifying AI-driven synchronization as a key research direction [3, 9]. Zhang et al. demonstrated Random Forest viability for real-time network anomaly detection with interpretable decision boundaries [10]. Prior work by Iliev and Penkov introduced the concept of AI-driven trust scores for synchronization control in 5G [11]. However, a common limitation in these studies is the predominant reliance on synthetic datasets, leaving the 'sim-to-real' generalization gap in operational transport networks largely unaddressed.

2.4. Positioning of AMSA.

AMSA differentiates itself in four aspects: (i) it addresses the complete management pipeline—from PDV filtering through source selection to holdover prediction; (ii) it employs an inherently explainable Random Forest whose feature importances and decision paths can be audited; (iii) it introduces contextual memory through a `prev_decision` feedback mechanism that prevents oscillatory switching; and (iv) it is validated against a hybrid dataset combining extensive field telemetry with critical corner-case simulations, ensuring robustness in real-world deployments.

3. AMSA architecture

AMSA is designed as a modular software overlay operating above the standard IEEE 1588 / ITU-T G.8275.1 protocol stack. It does not modify the underlying PTP message exchange but intercepts and augments the source selection decision. The architecture comprises four sequential stages (Fig. 1).




Fig. 1. System-level block diagram of the AMSA four-stage processing pipeline from raw PTP telemetry to adaptive source selection.

3.1. Stage 1: Initialization and Topological Analysis.

Upon deployment, AMSA performs topological discovery of all available timing references: GNSS receivers, PTP grandmasters via Announce messages, SyncE frequency inputs, and local oscillators. For each source, the system establishes a baseline statistical profile over $T_{init} = 300$ s, collecting mean and standard deviation of TE, MTIE at observation intervals $\tau = \{100, 1000, 10000\}$ s, one-way delay distribution parameters, and PTP Announce attributes.

3.2. Stage 2: Feature Engineering and Predictive Diagnostics.

1) PDV Filtering via Gamma Distribution: Standard PTP implementations apply linear filtering (PI controllers, Kalman filters) to measured delay, assuming Gaussian noise. In loaded IP/MPLS networks, PDV is consistently non-Gaussian: positive skewness (minimum delay is bounded; maximum is unbounded due to queuing) and heavy tails from cross-traffic bursts [5]. AMSA models one-way delay as $D(\alpha, \beta)$. Parameters are estimated via MLE on a rolling window of Sync/Delay_Resp timestamps. Packets exceeding the $(1 - \epsilon)$ quantile ($\epsilon = 0.01$) are rejected as outliers. Under high-PDV conditions ($\alpha = 0.5, \beta = 500$), this generates spikes reaching 2–3 μ s that defeat standard linear filters but are correctly identified and excluded by the Gamma model.

2) Virtual Rubidium — Holdover Prediction: AMSA continuously fits a polynomial regression (degree $k = 2 - 3$) to retrospective oscillator drift, computed from the difference between PTP-recovered time and the free-running clock over a sliding window. This captures second-order drift terms that BMCA's static holdover misses. The result is that TCXO oscillators (± 50 ppb, $\sim \text{€}100$) achieve holdover performance approaching miniature rubidium standards (± 0.01 ppb, $\sim \text{€}2,000$) through algorithmic compensation.

The complete feature vector comprises: instantaneous TE, rolling mean and standard deviation of TE (60 s window), MTIE, estimated Gamma parameters $\hat{\alpha}$ and $\hat{\beta}$, predicted holdover drift rate, and the previous arbitration decision (prev_decision) as a categorical variable.

3.3. Stage 3: Resilience Layer — Trust Management.

Per-source trust scores $T_i \in [0, 1]$ evolve based on statistical consistency with the weighted ensemble consensus. The trust for source i decays exponentially when its TE systematically diverges from the consensus mean. When T_i falls below the critical threshold $T_{crit} = 0.5$, the source is automatically demoted. This mechanism is effective against slow-drift GNSS spoofing, where a compromised receiver reports nominal clockClass 6 while its timing signal is gradually manipulated. The 250 s detection latency balances sensitivity against false positive rates during legitimate GNSS transients.

The spoofing detection capability represents a qualitative advance over BMCA. Standard PTP source selection evaluates only declared Announce attributes—clockClass, clockAccuracy, offsetScaledLogVariance—

none of which change during a well-executed spoofing attack. The manipulated receiver continues reporting clockClass 6 (GPS-locked), antenna status OK, and nominal variance. Only statistical analysis of the actual timing signal—specifically, the correlation between GNSS-reported phase and the PTP ensemble consensus—can detect the divergence. AMSA's trust mechanism performs exactly this analysis, providing a defense layer that no configuration of standard BMCA parameters can replicate.

3.4. Stage 4: Adaptive Arbitration (aBMCA).

The decision engine is a Random Forest classifier trained on a hybrid dataset combining extensive field telemetry and synthetic anomaly injection. To bridge the 'sim-to-real' gap, the model utilizes a massive repository of historical data from operational 5G nodes to learn nominal traffic patterns, while the simulation framework augments this dataset with labelled rare events (e.g., GNSS spoofing, extreme PDV) that are statistically scarce in real-world logs. Selection over deep learning was based on: inference speed (< 1 ms for real-time operation), robustness to feature scale differences, and explainability through feature importance (MDI) and decision path analysis [10]. The prev_decision input creates a first-order autoregressive loop: the model learns $P(y_t = A | x_t, y_{t-1} = A) > P(y_t = A | x_t, y_{t-1} = B)$, implementing adaptive hysteresis that varies with anomaly severity rather than requiring hard-coded thresholds.

4. Simulation methodology

AMSA is validated through a Monte Carlo framework (sync_sim.py) modeling the complete 5G timing chain. The training phase employs a Data Augmentation strategy: real-world PDV profiles derived from long-term network monitoring are injected with simulated fault scenarios (Classes B, C, D) to create a balanced, high-fidelity training set. This approach ensures the classifier remains robust against real-world noise while correctly identifying synthetic failure modes.

4.1. Signal Generation Models.

Oscillator Noise: AR (1) processes with five-level power-law noise following Allan deviation categorization (white PM, flicker PM, white FM, flicker FM, random walk FM). TCXO and OCXO profiles match manufacturer specifications (± 50 ppb and ± 0.01 ppb respectively). Discrete evolution is modeled as:

$$x(t) = x(t - 1) + d_{drift} + noise_{rw}. \quad (1)$$

PDV Model: Gamma-distributed $D(\alpha, \beta)$ with AR (1) temporal correlation ($\rho = 0.95$) modeling bursty congestion. Configurable path asymmetry (static and dynamic) and Bernoulli packet loss. Parameters are calibrated against field measurements.

Ground Truth Oracle: Accesses hidden error $E_i(t) = O_i(t) - T_{ref}(t)$. Labels are determined by:

$$Y(t) = \underset{i}{\operatorname{argmin}} |E_i(t)| + \lambda \cdot \sigma_i(t), \quad (2)$$

balancing accuracy vs. stability. Class balancing through systematic event injection yields $\sim 33\%/33\%/33\%$ distribution.

4.2. Evaluation Scenarios

Scenario A — Baseline: Low PDV, stable GNSS. Control experiment verifying zero AMSA overhead.

Scenario B — High PDV: Extreme congestion $D(0.5, 500)$ during $t = 1000 - 2000$ s, producing heavy-tailed spikes reaching $2-3 \mu\text{s}$.

Scenario C — GNSS Spoofing: Covert linear drift from $t = 1000$ s; GNSS reports nominal clockClass throughout.

Scenario D — GNSS Outage: Four phases: normal (0–500 s), GNSS loss (500–1500 s), full isolation (1500–2500 s), restoration (2500+ s).

Primary metrics: P95 TE, MTIE against G.8271.1 masks. Secondary: Max TE, mask violations, switch count.

5. Experimental results

5.1. Scenario A — Baseline Equivalence.

Both algorithms select GNSS with zero switches throughout the 3600 s window. Results are identical: P95 $TE = 432.15$ ns, Max $TE = 892.87$ ns, RMS $TE = 223.08$ ns, zero mask violations. The P95 of 432 ns provides substantial margin below the 1100 ns warning threshold. This identity—confirmed to the third decimal place—proves that AMSA introduces zero overhead, zero spurious switching, and zero additional noise when conditions are nominal. The algorithm is transparent: it adds complexity only when complexity is needed.

This result addresses a common concern regarding ML-based management systems: the risk of introducing unintended interference in stable conditions. Unlike some adaptive algorithms that produce spurious switching or computational noise even when no intervention is needed, AMSA demonstrates clean pass-through behavior. The algorithm's decision boundary is designed such that nominal conditions produce unanimous Random Forest consensus for the primary source, with prev_decision reinforcing this stability through learned inertia.

5.2. Scenario B — High-PDV Resilience.

Table 1. Performance Comparison Under High PDV (Scenario B).

Metric	Standard BMCA	AMSA	Δ
Max TE (ns)	1470.71	892.87	−39.3%
P95 TE (ns)	~980	446.51	−54.4%
RMS TE (ns)	—	223.08	—
Mask Violations	209	0	−100%
Switch Count	2	1,318	Adaptive

The introduction of extreme PDV reveals a pronounced performance gap. Standard BMCA permits PDV-corrupted packets through its servo, resulting in chaotic excursions reaching 1,471ns—dangerously close to the 1,500ns functional failure threshold. Over the test window, 209 mask violations are recorded.




Fig. 2. Comparative TE dynamics under Scenario B. BMCA (red) exhibits chaotic peaks breaching the $\pm 1,100$ ns mask boundary; AMSA (green) maintains stability throughout the high-PDV interval.

AMSA's Gamma filter correctly identifies and rejects heavy-tailed packets. Max $TE = 892.87$ ns—identical to baseline Scenario A—providing statistical evidence of complete compensation. P95 TE of 446.51 ns is only 14 ns above baseline, ensuring quality of service remains unaffected.




Fig. 3. Histogram of PDV in Scenario B, showing the heavy-tailed Gamma($\alpha=0.5$, $\beta=500$) distribution that defeats standard Gaussian-based filters.

The Stability Paradox. BMCA executed 2 switches; AMSA performed 1,318 (~ 1 per 2.7 s). Classical theory holds minimal switching is optimal. The data contradicts this: 1,318 switches produced zero violations; 2 switches produced 209. AMSA's switches are proactive: the classifier dynamically alternates between sources and holdover, constructing an optimal composite signal by gating out contaminated packets in real time. Each micro-switch isolates a PDV burst, then restores the primary source—a discrete optimization strategy that redefines stability as dynamic adaptivity rather than static persistence.

MTIE Spectral Analysis. The MTIE analysis provides spectral-domain confirmation. At medium observation intervals ($\tau = 10 - 100$ s)—the frequency range most affected by Gamma-distributed PDV—the BMCA curve approaches and crosses the G.8271.1 mask boundary, confirming that heavy-tailed congestion selectively corrupts mid-frequency stability. AMSA's MTIE curve shows smooth, bounded growth across all τ , proving that 1,318 micro-switches produce phase-coherent transitions rather than accumulative jitter. This evidence confirms that

adaptive micro-switching does not trade time-domain improvement for frequency-domain degradation.

5.3. Scenario C — GNSS Spoofing Detection.

A covert linear drift is injected at $t = 1000$ s. BMCA, evaluating declared attributes, tracks the manipulated signal—clockClass 6 remains nominal, antenna status reports “OK.” No standard PTP field indicates compromise.




Fig. 4. Trust score dynamics during spoofing. GNSS trust (blue) declines after attack onset; crosses $T_{crit}=0.5$ at $t \approx 1,250$ s. PTP trust (orange) remains stable. Lower panel: binary source selection showing the switching event.

AMSA’s trust mechanism detects statistical divergence between GNSS and the PTP ensemble consensus. The GNSS trust drops below $T_{crit,att} \approx 1,250$ s (250 s after attack onset), triggering automatic source switch. The TE stabilizes within ± 200 ns—well inside operational margins. Three characteristic phases are visible: accumulation (1000–1200 s, trust declining), crossing ($t \approx 1250$ s, trust < 0.5), and stabilization (PTP trust becomes dominant).

The 250 s detection latency is a design trade-off: faster detection increases false positive rates during legitimate GNSS transients; slower detection risks exceeding the timing budget. At the observed drift rate, 250 s produces ~ 200 – 300 ns accumulated error—within operational margin.

5.4. Scenario D — Extended Holdover and Recovery




Fig. 5. Comparative TE during spoofing. BMCA (red) follows injected drift; AMSA (green) isolates the compromised source and maintains stability through PTP reference.

The four-phase test evaluates: GNSS loss response ($t = 500$ s), deep holdover performance ($t = 1500 - 2500$ s), and recovery behavior ($t = 2500$ s).

Holdover Phase ($1500 - 2500$ s): Both systems rely on the local TCXO. BMCA applies a fixed frequency offset, leading to error accumulation as actual drift deviates from this static estimate—the TCXO’s second-order drift causes parabolic error growth. AMSA’s Virtual Rubidium continuously refines its polynomial drift model, predicting and compensating these terms.

Result: effective holdover duration (before $\pm 1.5 \mu s$) is extended by over 30%.




Fig. 6. Comparative TE through four phases of Scenario D. AMSA (green) maintains lower drift during holdover and achieves smooth recovery; BMCA (red) shows parabolic drift and recovery phase step.

Recovery Phase ($t = 2500$ s): Upon restoration, BMCA performs an abrupt frequency step correction, manifesting as a visible phase discontinuity—a transient mask violation. Quantitatively, BMCA’s instantaneous correction generates a phase step of approximately 800 ns—equivalent to broadband noise injection that produces a transient MTIE violation. AMSA implements a gradual frequency ramp over 60 s, achieving smooth convergence with maximum transient below 50ns—a $16\times$ reduction. This distinction is operationally critical: the recovery phase step can cause more service disruption than the holdover drift itself, as it occurs precisely when network services are being restored after an outage.

5.5. Field Validation

AMSA’s classification logic was applied to field data from an operational 5G transport network using Viavi MTS-5800 analyzers at multiple T-BC nodes [12]. Three scenarios were constructed:

Table 2. Field validation: expert vs. AMSA classification.

Source	Key Metric	Expert	AMSA (P)	Match
GNSS (failed)	TE=442513 ns	Critical	0.0001	✓
PTP L3 (ideal)	TE=2.3 ns	Reference	0.985	✓
Ethernet (boundary)	TE=100 ns	Marginal	0.72	✓

The GNSS failure case is noteworthy: TE of 442,513 ns exceeds the training range (0–5,000 ns) by $\sim 100\times$.

Despite this, Z-score normalization maps the value to $+500\sigma$, placing it unambiguously in the rejection zone of every decision tree. The boundary case (100 ns, exactly at PRTC-B limit) receives appropriate uncertainty ($P = 0.72$)—accepted but ranked below higher-confidence alternatives.

Table III consolidates all experimental results. Across four simulation scenarios and three field cases, AMSA demonstrates consistent superiority: zero-overhead in nominal conditions, complete PDV compensation, rapid spoofing detection, extended holdover with smooth recovery, and robust field generalization including out-of-distribution inputs.

5.6. Oscillator Sensitivity Analysis

Repeating Scenario B with TCXO (± 50 ppb) and Rubidium (± 0.05 ppb): AMSA+TCXO achieves P95 $TE \approx 440$ ns vs. BMCA+TCXO ≈ 980 ns (2.2× improvement). Critically, AMSA+TCXO outperforms BMCA + Rubidium (~ 500 ns), demonstrating that algorithmic compensation substitutes for hardware investment. AMSA + Rubidium yields only marginal gain (~ 430 ns) over AMSA + TCXO—the algorithm, not the oscillator, is the dominant factor.

The mechanism underlying this result is the Virtual Rubidium's polynomial drift model, which actively linearizes the TCXO error curve during holdover. While BMCA's static holdover allows second-order drift to accumulate quadratically (producing parabolic error growth), AMSA continuously predicts and compensates these terms, reducing the effective residual to approximately linear. This algorithmic compensation transforms inexpensive TCXO hardware into a virtual atomic reference—achieving the stability benefits of rubidium oscillators through software intelligence rather than hardware investment.




Fig. 7. P95 TE for all algorithm-oscillator combinations. AMSA+TCXO (€100) outperforms BMCA+Rubidium (€2,000).

6. Conclusion

This paper presented AMSA, a four-stage software overlay for resilient time source management in 5G networks. Experimental validation demonstrated: a 55% reduction in P95 Time Error; the elimination of 209 mask violations under high PDV conditions; robust spoofing detection within 250 s; and a 30%+ extension in holdover autonomy. Crucially, field validation confirmed 100% agreement with expert assessments, verifying the model's

capability to generalize beyond simulation data. The architecture achieves these gains without requiring hardware modifications or protocol changes to the ITU-T G.8275.1 stack.

The results highlight a "stability paradox": AMSA's 1,318 adaptive switches achieved zero violations, whereas the standard BMCA's 2 static switches resulted in 209 violations. This finding redefines operational stability as dynamic adaptivity rather than configuration persistence. Furthermore, the Virtual Rubidium concept proves that software intelligence can effectively substitute for expensive atomic oscillator hardware. By enabling standard TCXO components to outperform rubidium sources under AMSA control, the proposed system offers a pathway to significant infrastructure cost reduction.

Future work will extend this architecture to non-terrestrial network (NTN) synchronization using LEO satellite constellations, explore federated learning across multi-operator timing domains, and integrate with the O-RAN RIC for closed-loop synchronization management in 6G networks.

References

- [1] G. P. Sharma et al., "Toward Deterministic Communications in 6G Networks: State of the Art, Open Challenges and the Way Forward," *IEEE Access*, vol. 11, pp. 106898–106923, 2023.
- [2] A. Nasrallah et al., "Ultra-Low Latency (ULL) Networks: The IEEE TSN and IETF DetNet Standards and Related 5G ULL Research," *IEEE Commun. Surveys Tuts.*, vol. 21, no. 1, pp. 88–145, 2019.
- [3] D. Chandramouli, P. Andrés-Maldonado, and T. Kolding, "Evolution of Timing Services From 5G-A Toward 6G," *IEEE Access*, vol. 11, pp. 35150–35157, 2023.
- [4] IEEE, "IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems," *IEEE Std 1588-2019*, Jun. 2020.
- [5] M. Lévesque and D. Tipper, "A Survey of Clock Synchronization Over Packet-Switched Networks," *IEEE Commun. Surveys Tuts.*, vol. 18, no. 4, pp. 2926–2947, 2016.
- [6] F. Girela-López et al., "IEEE 1588 High Accuracy Default Profile: Applications and Challenges," *IEEE Access*, vol. 8, pp. 45211–45220, 2020.
- [7] G. Giorgi, "An Event-Based Kalman Filter for Clock Synchronization," *IEEE Trans. Instrum. Meas.*, vol. 64, no. 2, pp. 449–457, 2015.
- [8] Y.-D. Seo et al., "Improved Time-Synchronization Algorithm Based on Direct Compensation of Disturbance Effects," *Sensors*, vol. 19, no. 16, p. 3499, 2019.
- [9] A. Mahmood, "Industrial IoT in 5G-and-beyond networks: Vision, architecture, and design trends," *IEEE Trans. Ind. Informat.*, vol. 18, no. 6, pp. 4122–4137, Jun. 2022.
- [10] C. Zhang et al., "Three-Way Selection Random Forest Optimization Model for Anomaly Traffic Detection," *Electronics*, vol. 12, no. 8, p. 1788, 2023.
- [11] T. Iliev and I. Penkov, "AI-Driven Trust Score-Based Synchronization Control for Robust 5G and Beyond Networks," in *Proc. Int. Conf. Automatics and Informatics (ICAI)*, (in press).
- [12] P. Assimakopoulos et al., "A Converged Evolved Ethernet Fronthaul for the 5G Era," *IEEE J. Sel. Areas Commun.*, vol. 36, no. 11, pp. 2528–2537, 2018.