

Secure Near-Zero Latency Networking

OLATUNDE AYOMIDE OLASEHAN¹, UDOKA CYNTHIA DURUEMERUO², ADEBAYO ISHOLA³

¹*Computer Science, Swansea University, United Kingdom*

²*Computer Science, University of Wolverhampton, United Kingdom*

³*Computer Science, Swansea University, United Kingdom*

Abstract: The proliferation of real-time applications like autonomous vehicles and remote surgery has intensified the need for near-zero network latency with bounded jitter and deterministic behavior. However, latency optimizations through reduced protocol overhead and minimized buffering can expose attack surfaces, creating a security-latency trade-off where security measures violate timing constraints. This study examines approaches that co-design security with ultra-low-latency networking, focusing on edge-centric security, Zero Trust models, and secure transport protocols. Using exploratory analysis, the work evaluates architectures across latency impact, threat resilience, and scalability. The synthesis reveals patterns enabling secure near-zero latency: edge-based processing with localized trust anchors, hardware-accelerated encryption, and pre-established secure sessions. Low-overhead techniques include deterministic authentication and selective enforcement based on traffic criticality. Findings show centralized security is incompatible with near-zero latency, while distributed enforcement can maintain microsecond-scale security overhead. An integrated model combining deterministic transport and localized trust demonstrates that secure, ultra-low-latency networking is achievable through co-designed architectures.

I. INTRODUCTION

The rapid evolution of digital infrastructures has led to a growing class of applications that demand near-zero end-to-end latency, often measured in microseconds to a few milliseconds. Real-time and mission-critical systems such as autonomous vehicles, industrial automation, remote and robotic surgery, high-frequency trading, and immersive extended reality (XR) environments rely on immediate feedback and deterministic network behavior to function correctly and safely (Popovski et al., 2019). Even minor latency variations in such systems can lead to degraded performance, instability, or catastrophic failure.

To meet these requirements, modern networking paradigms increasingly employ edge computing,

Time-Sensitive Networking (TSN), Ultra-Reliable Low-Latency Communications (URLLC) in 5G, kernel-bypass techniques such as RDMA, and advanced congestion-control mechanisms that reduce queuing delays. These developments reflect a fundamental shift away from best-effort networking toward deterministic, latency-bounded architectures.

However, as networks become faster and more distributed, they also become more exposed to security threats, including denial-of-service attacks, man-in-the-middle interception, data manipulation, and unauthorized access. Ensuring security in such environments without compromising stringent latency requirements remains a central challenge in next-generation network design.

Achieving near-zero latency typically requires aggressive optimization techniques, such as minimizing protocol overhead, bypassing kernel processing, reducing buffering, and shortening or eliminating connection setup phases. While these optimizations improve responsiveness, they often weaken or bypass traditional security mechanisms, creating new attack surfaces (Briscoe et al., 2016).

Conventional network security architectures such as deep packet inspection, centralized authentication services, and multi-round cryptographic handshakes introduce non-negligible processing and communication delays that conflict directly with ultra-low latency objectives. For example, standard TLS handshakes, intrusion detection systems, and centralized policy enforcement can add milliseconds of delay, which is unacceptable in time-critical systems like industrial control loops or remote medical procedures (Rescorla, 2018).

This tension creates a fundamental performance-security trade-off, where strengthening security can

degrade latency, while prioritizing speed can reduce protection. Existing solutions often address latency and security in isolation, leading to fragmented architectures that fail to provide holistic guarantees for both dimensions.

The purpose of this study is to systematically examine architectural approaches that enable secure near-zero latency networking while maintaining acceptable trade-offs among performance, reliability, and security guarantees. Rather than treating security as an add-on, this work focuses on co-designed architectures where security mechanisms are integrated into low-latency networking stacks.

Specifically, the study analyzes emerging paradigms such as edge-centric security, Zero Trust networking models, lightweight cryptography, secure transport protocols (e.g., QUIC and TLS 1.3), and cross-layer optimization techniques that jointly consider latency and threat resilience (Rescorla, 2018). By examining these approaches, the study aims to identify design patterns and architectural principles suitable for real-time and mission-critical environments.

This study contributes to both academic research and practical network engineering by addressing a problem of growing importance across multiple industries. As societies increasingly rely on real-time digital systems, failures caused by either excessive latency or inadequate security can have severe economic, safety, and societal consequences.

The findings provide design guidance for engineers developing ultra-low latency networks, assist researchers in identifying open challenges and trade-offs, and inform standards bodies involved in shaping future networking technologies such as 5G/6G, TSN, and next-generation transport protocols. Ultimately, this work supports the development of secure, resilient, and deterministic networking infrastructures capable of supporting the next wave of real-time and mission-critical applications.

II. LITERATURE REVIEW

2.1 Near-Zero Latency Networking Concepts

Near-zero latency networking refers to communication systems engineered to achieve extremely low end-to-end delays, typically ranging from microseconds to a few milliseconds, with tightly bounded jitter. While “zero latency” is theoretically unattainable, the term is commonly used to describe networks designed to meet strict latency budgets imposed by real-time and mission-critical applications such as industrial control, autonomous systems, and tactile Internet services (Popovski et al., 2019). Latency budgets in such systems account for transmission, propagation, processing, and queuing delays, each of which must be minimized or tightly controlled.

Time-Sensitive Networking (TSN) is a key enabling technology for near-zero latency communication in wired environments. TSN, standardized by IEEE 802.1, introduces mechanisms such as time-aware shaping, traffic scheduling, and precise clock synchronization to provide deterministic latency and bounded jitter over Ethernet networks. These features make TSN suitable for industrial automation, automotive networks, and real-time control systems.

Edge computing complements TSN by reducing latency through the relocation of computation and storage closer to data sources. By minimizing round-trip delays to centralized cloud infrastructures, edge architectures enable rapid processing and response, which is essential for applications with stringent real-time constraints (Satyanarayanan, 2017).

Software-Defined Networking (SDN) further contributes to near-zero latency by enabling centralized and programmable control over network behavior. SDN allows dynamic traffic prioritization, optimized path selection, and rapid reconfiguration in response to congestion or failures, supporting latency-aware traffic engineering. Together, TSN, edge computing, and SDN form the technological foundation of modern near-zero latency networking systems.

2.2 Security Challenges in Ultra-Low Latency Networks

Despite advances in latency reduction, incorporating security into ultra-low latency networks remains

challenging. Traditional security mechanismssuch as encryption, authentication, and integrity verification, introduce computational and communication overhead that can conflict with strict latency budgets (Rescorla, 2018). In real-time environments, even small processing delays caused by cryptographic operations or multi-round authentication handshakes may violate application-level timing constraints.

Encryption ensures confidentiality but adds processing overhead and increases packet sizes, which can affect transmission time. Authentication mechanisms, particularly those relying on public-key cryptography or centralized identity management, introduce handshake delays that are problematic for latency-critical systems. Similarly, integrity checks and deep packet inspection improve resilience against data manipulation and attacks but often require additional computation or buffering, increasing end-to-end latency.

The literature consistently highlights a security-latency trade-off, where strengthening security controls can degrade performance, while aggressive latency optimization may weaken protection against threats such as denial-of-service attacks, spoofing, or unauthorized access (Briscoe et al., 2016). This trade-off is particularly pronounced in distributed and edge-based architectures, where security enforcement must be performed close to the data plane without centralized oversight.

2.3 Existing Low-Latency Network Architectures

Several architectural paradigms have been proposed to support ultra-low latency communication. Among these, 5G Ultra-Reliable Low-Latency Communication (URLLC) is one of the most prominent. URLLC, standardized by the 3GPP, targets end-to-end latencies below 1 ms with extremely high reliability, enabling applications such as industrial automation, vehicle-to-everything (V2X) communication, and remote control systems (Popovski et al., 2019). URLLC achieves these targets through mechanisms such as shortened transmission time intervals, prioritized scheduling, and optimized radio resource allocation.

Edge-centric architectures extend low-latency guarantees beyond the radio access network by placing computation and security functions at edge nodes. This architectural approach reduces dependency on centralized cloud services and improves responsiveness for time-critical workloads (Satyanarayanan, 2017).

Deterministic networking, particularly TSN, provides predictable latency and bounded jitter in wired networks. Research has increasingly explored the integration of deterministic networking with wireless technologies such as 5G to enable end-to-end latency guarantees across heterogeneous infrastructures (Finn et al., 2019).

Additionally, in-network computing has emerged as an approach for reducing latency by performing computation directly within the network fabric. By offloading tasks such as aggregation, filtering, or telemetry to programmable switches and network devices, in-network computing can significantly reduce processing delays and bandwidth usage (Sapiro et al., 2017).

2.4 Security Mechanisms for Low-Latency Environments

To address the security-latency trade-off, researchers have proposed security mechanisms specifically optimized for low-latency environments. Lightweight cryptography is designed to provide confidentiality and integrity with reduced computational complexity, making it suitable for real-time and resource-constrained systems (Beaulieu et al., 2015).

Hardware-assisted security leverages cryptographic accelerators and trusted execution environments to offload security operations from software, significantly reducing processing delays. Technologies such as AES-NI and FPGA-based security modules enable high-throughput encryption with minimal latency overhead.

Pre-authentication and trust anchoring techniques aim to eliminate or minimize runtime authentication overhead by establishing trust relationships prior to data exchange. Pre-shared keys, cached credentials, and localized trust anchors reduce the need for

repeated handshakes, thereby improving responsiveness in latency-sensitive scenarios (Iyengar & Thomson, 2021).

Furthermore, physical layer security (PLS) has been explored as an alternative or complement to traditional cryptography. PLS exploits the physical characteristics of wireless channels to provide confidentiality and resistance to eavesdropping with minimal processing overhead. While promising, PLS techniques often require integration with higher-layer security mechanisms to ensure comprehensive protection.

2.5 Identified Gaps in Current Literature

Despite extensive research on low-latency networking and security optimization, several gaps remain. First, most studies address latency and security in isolation, with limited holistic analyses that jointly evaluate their interaction and trade-offs in real-world systems. Second, there is a lack of unified architectural frameworks that integrate deterministic networking, edge computing, and security mechanisms under common design principles.

Additionally, standardized evaluation methodologies that simultaneously measure latency, reliability, and security resilience are scarce, making it difficult to compare proposed solutions across studies. These gaps highlight the need for comprehensive frameworks and cross-layer approaches that treat security as a core component of near-zero latency network design rather than an afterthought.

III. THEORETICAL FRAMEWORK

3.1 Latency–Security Trade-off Model

Modern real-time networks must simultaneously satisfy two often-competing objectives: ultra-low end-to-end latency and robust security guarantees. In this context, a Latency–Security Trade-off Model conceptualizes how security controls influence processing overhead and overall network latency.

At a high level, total end-to-end latency (L_{E2E}) in a communication system can be viewed as a sum of constituent components such as transmission,

propagation, queuing, processing, and security overhead:

$$L_{E2E} = L_{\text{transmission}} + L_{\text{propagation}} + L_{\text{queuing}} + L_{\text{processing}} + L_{\text{security}}$$

Here, represents additional time contributed by cryptographic computation, handshakes, authentication, and integrity verification. As the strength and complexity of security controls increase, generally increases, leading to a trade-off between security resilience and latency performance. For example, cryptographically strong authentication and multi-round handshakes increase assurance against unauthorized access but introduce delays that are unacceptable in ultra-reliable low-latency communication (URLLC) contexts (Gallenmüller et al., 2020).

This conceptual model captures the optimization challenge faced in secure near-zero latency networks: minimizing latency while maintaining sufficient security strength. Security mechanisms that reduce overhead — such as pre-authentication, hardware acceleration, or lightweight cryptographic primitives — effectively shift the trade-off curve, enabling stronger security at the same latency budget or similar security with lower latency cost.

Graphically, such a model is often represented as two opposing curves on a performance plane: latency on the horizontal axis and security strength on the vertical axis. A steeper security curve represents high cost to latency performance as security increases; flatter curves denote optimized security mechanisms with lower latency penalty. Practically, the design space is constrained by application requirements, such as the millisecond-scale bounds needed for industrial control or autonomous vehicle communications. This framework helps guide system designers toward balance points where latency requirements and security guarantees coexist within acceptable risk thresholds.

3.2 Trust and Determinism in Real-Time Networks

Real-time network systems, particularly those supporting mission-critical applications, depend not only on low latency but also on trust establishment and

deterministic performance guarantees. Unlike traditional packet networks where delivery times are probabilistic, deterministic networking aims to ensure bounded latency and low jitter for admitted flows, often through scheduling, resource reservation, and synchronized time enforcement.

Trust establishment refers to mechanisms that ensure participating entities are authenticated, authorized, and operating under agreed security policies before engaging in deterministic communication flows. In the context of near-zero latency networking, trust must often be established with minimal handshake overhead, necessitating techniques such as pre-shared credentials, cached trust associations, or local trust anchors that reduce runtime verification delays. The result is a model where trust membership is established once or infrequently, and then secure deterministic data exchange can proceed rapidly without repeated heavy security negotiations.

Deterministic communication, formalized in standards such as Time-Sensitive Networking (TSN) and IETF Deterministic Networking (DetNet), supports predictable performance through traffic shaping, scheduling, and precise timing synchronization. These mechanisms guarantee that traffic flows meet strict delay bounds by allocating dedicated network resources and avoiding traditional congestion variability. Deterministic frameworks make it possible to deliver data with minimal jitter and bounded latency, an essential requirement for applications like industrial automation and remote robotic control.

Integrating trust into deterministic networking enhances predictable performance by ensuring that only authenticated and authorized flows can reserve deterministic resources. Such integration is crucial because unauthorized or malicious flows could consume critical bandwidth or disrupt scheduled paths, undermining performance guarantees. A theoretical framework for trust and determinism thus includes:

Trust provisioning layer: Establishes identity and authorization with minimal latency impact, often through hybrid cryptographic and hardware trust anchors.

Resource reservation and scheduling plane: Allocates deterministic slots or reserved resources to trusted flows.

Performance enforcement layer: Monitors and ensures compliance with latency and jitter guarantees.

By fusing trust establishment with deterministic scheduling, networks can reliably support near-zero latency communication while resisting unauthorized access and performance degradation — a necessary design principle for real-time, mission-critical environments.

IV. METHODOLOGY

4.1 Research Design

This study adopts an exploratory qualitative research design to examine architectural and security approaches for secure near-zero latency networking. Exploratory research is particularly suited to investigations where existing knowledge is fragmented and formal models remain underdeveloped (Stebbins, 2001). In the context of next-generation low-latency networks, many proposed solutions span disparate domains—networking standards, edge computing paradigms, hardware acceleration techniques, and lightweight cryptographic schemes—without consolidated evaluation frameworks (Briscoe et al., 2016).

The research design incorporates architectural comparison and synthesis of existing studies to identify latent design patterns, trade-offs, and gaps. This involves systematically reviewing relevant literature to abstract key characteristics of each architectural approach and security mechanism, then synthesizing insights to inform high-level design principles. Such synthesis supports constructive alignment across latency impact, security resilience, and scalability, facilitating theoretical generalization while maintaining relevance to real-world systems (Okoli & Schabram, 2011).

4.2 Data Sources

Primary data sources include:

Peer-reviewed journals in networking, security, and communications systems (e.g., IEEE Communications

Surveys & Tutorials, ACM Computing Surveys, IEEE Internet of Things Journal), which provide rigorous empirical and conceptual analyses of latency mechanisms and security trade-offs (Popovski et al., 2019).

Industry white papers and technical reports published by standards bodies and consortia (e.g., IEEE 802.1 TSN Task Group, 3GPP URLLC specifications), which offer up-to-date definitions, architectural models, and performance targets relevant to deterministic and low-latency networking.

Standards documentation and Internet Engineering Task Force (IETF) Request for Comments (RFCs) that define protocols and architectural frameworks (e.g., TLS 1.3, QUIC) affecting secure communication performance (Iyengar & Thomson, 2021).

These sources ensure methodological rigor by grounding analysis in established academic discourse and industrial practice, capturing both theoretical proposals and real-world implementation considerations.

4.3 Analysis Technique

The analytical strategy consists of a comparative analysis framework that assesses network architectures and security mechanisms along three core dimensions:

Latency Impact: Evaluation of how specific architectural choices and security controls affect end-to-end latency, including overhead introduced by processing, handshakes, and communication rounds.

Resilience: Assessment of each mechanism's ability to withstand security threats such as unauthorized access, message tampering, replay attacks, and denial-of-service without degrading performance beyond acceptable bounds.

Scalability: Consideration of whether solutions can maintain performance and security guarantees as network size, node heterogeneity, and traffic loads increase, a critical concern for edge and distributed infrastructures.

The comparative analysis is executed through structured tabulation and narrative synthesis, enabling identification of commonalities and divergences among candidate solutions. For example, lightweight cryptography mechanisms are compared against hardware-assisted approaches in terms of latency overhead, resource requirements, and applicability to edge nodes. Similarly, 5G URLLC, TSN, and SDN-enabled architectures are evaluated based on latency budgeting strategies and integration of security controls.

This approach provides a multi-faceted evaluation that balances descriptive richness with analytical clarity, revealing not only what architectural options exist but also how they trade off performance and security in near-zero latency environments.

V. RESULTS

5.1 Key Architectural Patterns for Secure Near-Zero Latency

The analysis of existing architectures and security mechanisms reveals a set of recurring architectural patterns that enable secure communication while preserving near-zero latency guarantees. Edge-based processing and localized trust is a foundational pattern across low-latency systems. By relocating computation, policy enforcement, and trust management to edge nodes, networks significantly reduce round-trip delays to centralized cloud services. Localized trust anchorssuch as pre-provisioned credentials, local certificate authorities, or Zero Trust edge gateways—enable rapid authentication and authorization without repeated wide-area signaling. This pattern aligns with edge computing paradigms that emphasize autonomy, responsiveness, and context-aware decision-making in real-time environments(Satyanarayanan, 2017).

Hardware-accelerated encryption and inspection further emerges as a critical enabler of secure low-latency networking. Cryptographic accelerators, SmartNICs, and trusted execution environments allow encryption, decryption, and integrity checks to be performed at line rate, reducing processing overhead compared to software-based approaches. Hardware-assisted security is particularly effective in

deterministic networks, where predictable execution times are required to maintain bounded latency and low jitter.

Pre-established secure sessions represent another key pattern. Rather than performing costly authentication and key exchange during time-critical operation, trust relationships are established in advance through pre-shared keys, cached credentials, or long-lived secure associations. This approach minimizes handshake delays and ensures that secure data transmission can begin immediately when real-time constraints are most stringent (Iyengar & Thomson, 2021).

Together, these architectural patterns demonstrate that security mechanisms must be embedded within low-latency architectures rather than added as reactive layers.

5.2 Security Techniques with Minimal Latency Overhead

Beyond architectural patterns, the findings highlight specific security techniques that achieve protection goals with minimal impact on latency.

Lightweight cryptographic protocols are widely adopted in latency-sensitive environments due to their reduced computational complexity. Optimized symmetric encryption and streamlined integrity mechanisms provide essential confidentiality and authenticity while respecting strict latency budgets, particularly in embedded and industrial systems (Beaulieu et al., 2015).

Deterministic authentication mechanisms replace probabilistic, multi-round handshakes with bounded and predictable authentication processes. Techniques such as pre-authentication, credential caching, and identity pre-provisioning ensure that authentication latency remains consistent and does not introduce jitter or unexpected delays. This predictability is essential for real-time applications operating under deterministic scheduling constraints (Rescorla, 2018).

Selective and adaptive security enforcement enables security controls to be applied contextually rather than uniformly. By tailoring security intensity to traffic criticality, trust level, and operational context, networks can protect high-risk flows without imposing

unnecessary overhead on latency-critical traffic. This adaptive approach supports fine-grained balancing between performance and security objectives in heterogeneous edge environments (Briscoe et al., 2016).

5.3 Performance and Security Trade-offs

The results reaffirm that security inevitably introduces overhead, but the magnitude and variability of this overhead depend strongly on architectural and design choices. The impact of security layers on latency budgets varies significantly. Centralized authentication services, deep packet inspection, and multi-round cryptographic handshakes introduce delays that are incompatible with near-zero latency requirements. In contrast, distributed enforcement, hardware acceleration, and pre-established trust relationships substantially reduce added latency, often keeping security overhead within acceptable microsecond-scale bounds.

Balancing resilience, fault tolerance, and real-time guarantees presents an additional trade-off. Redundancy and fault tolerance improve availability but may introduce extra signaling, synchronization, or failover delays. The findings indicate that resilience mechanisms must be tightly integrated with deterministic scheduling and local control to avoid violating real-time guarantees during failure conditions.

Overall, the trade-off analysis shows that secure near-zero latency networking is achievable when security mechanisms are co-designed with performance objectives and deterministic behavior.

5.4 Integrated Secure Near-Zero Latency Networking Model

Based on the synthesized findings, this study proposes a high-level Integrated Secure Near-Zero Latency Networking Model that unifies deterministic networking with low-overhead security controls.

The model consists of three tightly coupled layers:

1. Deterministic Transport and Scheduling Layer

Provides bounded latency and minimal jitter using TSN, DetNet, URLLC, and latency-aware traffic engineering. Network resources are explicitly reserved and scheduled to eliminate congestion variability (Finn et al., 2019).

2. Low-Overhead Security Enforcement Layer

Implements lightweight cryptography, hardware-accelerated encryption, and selective inspection directly within the data plane, ensuring security operations execute within predictable timing bounds.

3. Localized Trust and Control Layer

Establishes trust through pre-provisioned identities, local trust anchors, and infrequent authentication events, enabling rapid and secure flow admission without repeated handshake delays.

By integrating these layers, the model supports predictable performance while maintaining robust security guarantees, offering a coherent architectural foundation for real-time and mission-critical systems.

VI. CONCLUSION

This study demonstrates that secure near-zero latency networking is achievable through careful architectural and security design choices. The central finding is that latency and security are not inherently opposing objectives when security mechanisms are embedded within deterministic, edge-centric architectures. Hardware-assisted security, pre-established trust, lightweight cryptography, and selective enforcement significantly reduce latency overhead while preserving essential protection against threats.

Rather than relying on best-effort security models, future real-time networks must adopt co-designed approaches that jointly optimize performance, reliability, and security. Such integration is essential for supporting emerging applications that depend on deterministic behavior and rapid response.

VII. RECOMMENDATIONS

For Network Architects and System Designers Integrate security mechanisms directly into low-

latency architectures rather than treating them as external layers. Leverage edge computing, hardware acceleration, and deterministic scheduling to minimize security-related delays. Explicitly account for security overhead within latency budgets during system design. For Developers of Real-Time and Mission-Critical Applications Employ pre-authentication and session reuse to avoid runtime handshake delays. Use lightweight and hardware-accelerated cryptographic libraries optimized for real-time execution. Design applications to align with deterministic networking assumptions and constraints. For Standards Bodies and Policy Stakeholders Promote unified frameworks that jointly address latency, security, and determinism. Encourage standardization of low-overhead security mechanisms for TSN, 5G/6G, and edge environments. Support evaluation methodologies that measure latency, reliability, and security resilience together.

REFERENCES

- [1] Beaulieu, R., Shors, D., Smith, J., Treatman-Clark, S., Weeks, B., & Wingers, L. (2015). The Simon and Speck Block Ciphers on AVR 8-Bit Microcontrollers (pp. 3–20). https://doi.org/10.1007/978-3-319-16363-5_1
- [2] Briscoe, B., Brunstrom, A., Petlund, A., Hayes, D., Ros, D., Tsang, I.-J., Gjessing, S., Fairhurst, G., Griwodz, C., & Welzl, M. (2016). Reducing Internet Latency: A Survey of Techniques and Their Merits. *IEEE Communications Surveys & Tutorials*, 18(3), 2149–2196. <https://doi.org/10.1109/comst.2014.2375213>
- [3] Finn, N., Thubert, P., Varga, B., & Farkas, J. (2019). Deterministic Networking Architecture (Vol. 8655). <https://doi.org/10.17487/rfc8655>
- [4] Gallenmüller, S., Naab, J., Adam, I., & Carle, G. (2020). 5G QoS: Impact of Security Functions on Latency. 1–9. <https://doi.org/10.1109/noms47738.2020.9110422>
- [5] Iyengar, J., & Thomson, M. (2021). QUIC: A UDP-Based Multiplexed and Secure Transport. <https://doi.org/10.17487/rfc9000>
- [6] Okoli, C., & Schabram, K. (2011). A Guide to Conducting a Systematic Literature Review of Information Systems Research. *SSRN Electronic Journal*, 10(26). <https://doi.org/10.2139/ssrn.1954824>

- [7] Popovski, P., Stefanovic, C., Nielsen, J. J., De Carvalho, E., Angelichinoski, M., Trillingsgaard, K. F., & Bana, A.-S. (2019). Wireless Access in Ultra-Reliable Low-Latency Communication (URLLC). *IEEE Transactions on Communications*, 67(8), 5783–5801.
<https://doi.org/10.1109/tcomm.2019.2914652>
- [8] Rescorla, E. (2018). The Transport Layer Security (TLS) Protocol Version 1.3 (Vol. 8446). <https://doi.org/10.17487/rfc8446>
- [9] Sapiro, A., Abdelaziz, I., Aldilaijan, A., Canini, M., & Kalnis, P. (2017). In-Network Computation is a Dumb Idea Whose Time Has Come. 150–156.
<https://doi.org/10.1145/3152434.3152461>
- [10] Satyanarayanan, M. (2017). The Emergence of Edge Computing. *Computer*, 50(1), 30–39.
<https://doi.org/10.1109/mc.2017.9>
- [11] Stebbins, R. (2001). Exploratory Research in the Social Sciences. Sage.
<https://doi.org/10.4135/9781412984249>