



I. D. Ivanets¹, V. K. Ovsyak^{1,2}

¹ Lviv Polytechnic National University, Lviv, Ukraine

² Kielce University of Technology, Kielce, Poland

EDGE-NATIVE CABLE ACCESS NETWORK WITH QUIC TRANSPORT OVER UDP

This paper introduces a novel architectural model for cable access networks that leverages an edge-native approach to relocate QUIC protocol termination from centralized cores to the access edge. The proposed model addresses the performance limitations of traditional DOCSIS-based architectures, particularly under asymmetric traffic conditions and real-time service demands. To evaluate the model viability, a series of controlled experiments were conducted using the FD.io Vector Packet Processing (VPP) framework, the vpecho benchmarking tool, and the CSIT continuous performance testing infrastructure. Tests were executed on 3-node ICX-100GE testbed with Intel Ice Lake processors and 100GE Intel E810CQ interfaces, using a fixed network frame size of 1280 bytes and a single CPU core per instance. Four QUIC-over-UDP scenarios were analyzed: a baseline configuration (1 client, 1 session) achieving 882.8 Mbps; a 10-client, single-session case with 625.3 Mbps; a single-client, 10-session case with 586.5 Mbps; and a 10-client, 10-session configuration yielding 562 Mbps. The results show that the baseline configuration offers the highest throughput, while scaled session and client loads introduce measurable performance degradation due to increased overhead. This degradation is attributed to session tracking overhead, NUMA effects, and per-client context switching in a constrained single-threaded environment, highlighting the importance of scaling strategies. Its flexibility and software-defined nature make it applicable to evolving requirements in low-latency and scalable broadband services. These findings validate the efficiency and scalability of user-space QUIC transport at the network edge. This work lays the foundation for further research into edge-hosted QUIC processing in distributed DOCSIS environments. The created edge-native model demonstrates its suitability for cost-effective stepwise modernization of optical-coaxial access networks, offering improved responsiveness and reduced reliance on centralized infrastructure without requiring physical plant replacement.

Keywords: transport-layer offloading; asymmetrical broadband; high-performance packet processing; scalable configuration; network softwarization; performance benchmarking.

Introduction / Вступ

The exponential growth of upstream-demanding services – such as interactive video conferencing, cloud-based collaboration, and telemedicine – has placed increasing pressure on traditional broadband architectures. Among these, hybrid fiber-coaxial (HFC) networks, widely deployed across North America and Europe, face growing limitations due to their inherent asymmetry and centralized transport-layer processing. Although standards like DOCSIS 3.1/4.0 and architectural extensions such as Remote PHY Devices (RPD) and Distributed Access Architectures (DAA) aim to extend capacity and fiber reach, they fall short of addressing key transport-layer bottlenecks in asymmetric environments [7]. This challenge has been addressed in mobile networking through Multi-access Edge Computing (MEC) and software-defined networking (SDN), which have enabled low-latency services via functional relocation to the access edge [5]. However, similar approaches in fixed broadband – particularly HFC – remain insufficiently explored.

A growing number of studies highlight the potential of softwarized user-space implementations of transport proto-

cols as a scalable solution to these challenges. One such protocol is QUIC, developed by Google and standardized by the IETF, which integrates multiplexed streams, TLS 1.3, and congestion control over UDP in a fully user-space model [4]. While QUIC has been successfully deployed in data center and mobile ecosystems, its practical performance under asymmetric conditions – typical of HFC networks – has not yet been experimentally validated in access-layer deployments. Moreover, publicly available benchmarking in the FD.io VPP ecosystem lacks systematic evaluation of QUIC transport termination on softwarized, non-centralized network nodes.

This study proposes a new architectural model for cable access networks that enables QUIC-over-UDP termination to be deployed at the aggregation layer, close to the user. Unlike centralized virtual CMTS (vCMTS) models, this approach leverages user-space packet processing via Vector Packet Processing (VPP) and Continuous System Integration Testing (CSIT) to deliver measurable performance benefits while avoiding physical infrastructure upgrades. Benchmarking is carried out on the ICX-3N testbed consisting of three Intel Ice Lake nodes with 100GE Intel

Інформація про авторів:

Іванець Іван Дмитрович, аспірант, кафедра комп'ютерних технологій у видавничо-поліграфічних процесах.

Email: ivan.d.ivanets@lpnu.ua; <https://orcid.org/0009-0004-7508-4867>

Овсяк Володимир Казимирович, д-р техн. наук, професор, кафедра комп'ютерних технологій у видавничо-поліграфічних процесах. Email: vovsyak@tu.kielce.pl; <https://orcid.org/0000-0001-9295-284X>

Цитування за ДСТУ: Іванець І. Д., Овсяк В. К. Edge-native архітектура гібридної оптично-коаксальної мережі з реалізацією Quic-протоколу через UDP. Науковий вісник НЛТУ України. 2025, т. 35, № 3. С. 114–119.

Citation APA: Ivanets, I. D., & Ovsyak, V. K. (2025). Edge-native cable access network with Quic transport over UDP. *Scientific Bulletin of UNFU*, 35(3), 114–119. <https://doi.org/10.36930/40350312>

E810CQ NICs, using the vpecho test suite, a fixed network frame size of 1280 bytes, and a single CPU core per test process. The proposed model is validated through performance measurements in four QUIC configurations: baseline (882.8 Mbps), 10 clients with one session each (625.3 Mbps), one client with 10 sessions (586.5 Mbps), and 10 clients with 10 sessions (562 Mbps). These results confirm that the highest throughput is achieved under the baseline configuration, while session and client scaling introduces performance degradation due to increased transport overhead. The approach demonstrates consistent results across scenarios and shows potential for phased HFC modernization through transport-layer decentralization.

Object of research – building a hybrid fiber-coaxial network architecture.

Subject of research – methods and means of building an architecture for HFC access networks with QUIC protocol implementation over UDP, which enables improved upstream throughput, reduced latency, and software-defined scalability without physical infrastructure replacement.

The purpose of the research – to develop and evaluate an edge-native architecture for HFC networks with QUIC transport over UDP, which enables transport-layer functionality to be relocated to the access edge without requiring full physical infrastructure replacement.

To achieve this *purpose*, the following main *research objectives are identified*:

- to analyze the latest research and publications, in order to outline the current scientific context, identify knowledge gaps, and justify the relevance of the proposed approach;
- to create a model of the HFC network architecture, which enables transport-layer functionality to be relocated to the access edge without requiring full physical infrastructure replacement;
- to experimentally compare baseline and multi-session transport termination scenarios using the vpecho benchmarking tool within the CSIT automation framework, in order to assess the real-world performance of user-space QUIC implementations.

Analysis of recent research and publications. In the report [15], the limitations of DOCSIS 3.1-based HFC networks are clearly outlined, particularly the upstream constraint of 35 Mbps due to the 42 MHz spectral split and coaxial return path. The document emphasizes that, while downstream capacity can reach up to 1 Gbps, the limited upstream bandwidth significantly impairs performance under modern use cases such as teleconferencing and real-time collaboration. This evidence highlights the structural asymmetry inherent to DOCSIS 3.1, thereby motivating the need for architectural changes aimed at enhancing upstream transport efficiency.

According to the CableLabs DOCSIS 3.1 Physical Layer Specification [16], HFC networks are limited in upstream capacity due to a fixed spectrum allocation of 5-85 MHz for return path traffic, while downstream operates over a broader 108-860 MHz range. This asymmetrical bandwidth allocation was originally suitable for downstream-heavy services, but it no longer aligns with the requirements of modern applications, which demand higher upstream performance and low-latency bidirectional communication.

The DOCSIS 4.0 Physical Layer Specification published by CableLabs [12] introduces two key enhancements aimed at overcoming the upstream bandwidth limitations of traditional HFC networks: Extended Spectrum DOCSIS (ESD) and Full Duplex DOCSIS (FDX). The ESD mode

expands the usable coaxial spectrum up to 1.8 GHz, while the FDX mode allows for simultaneous upstream and downstream transmissions over overlapping frequency ranges. According to the specification, the deployment of FDX technology imposes significant operational challenges, including the need for echo cancellation, plant balancing, and stringent tolerance control in the outside plant. These factors complicate the real-world adoption of DOCSIS 4.0.

In the study [6], the authors emphasize that although the global number of internet users increased by 1.5 billion between 2018 and 2022, the benefits of digitalization remain unevenly distributed. Fixed broadband penetration in low-income countries remains close to zero, and only one in four people in those regions had internet access in 2022. The report highlights that affordability, poor infrastructure, and limited digital literacy are the main barriers to adoption. These digital divides are especially acute in suburban and rural areas, where the transition to high-quality symmetrical broadband remains slow. This reinforces the need to explore alternative modernization models for cable access networks that avoid full infrastructure replacement and remain cost-effective in underserved areas.

In their study [7], the authors examined the operational behavior of large-scale RPD deployments in DAA. They highlighted that while the migration of physical layer processing to the access node improves signal quality and simplifies maintenance, the centralized nature of MAC and IP-layer functions in vCMTS environments poses serious scalability limitations. The authors reported that under high concurrency – particularly when handling thousands of simultaneous subscriber flows – the vCMTS architecture suffers from latency spikes and resource contention. This reinforces the need to explore edge-localized alternatives that can address transport-layer congestion more effectively.

In the comprehensive survey [5] reviewed the evolution and challenges of SDN, highlighting its potential to decouple the control and data planes, improve network programmability, and enable fine-grained traffic steering in both cloud and mobile environments. The study emphasizes that SDN, in conjunction with network function virtualization (NFV), lays the foundation for flexible service deployment at the network edge. These findings support the exploration of programmable user-space transport technologies in edge-constrained access networks such as HFC, where traditional transport architectures face scalability and latency bottlenecks.

In the article [1] presented the FD.io VPP framework as a high-performance, user-space network stack optimized for modern multi-core x86 architectures. The authors demonstrated that VPP leverages vectorized processing and a polling-based model to achieve terabit-scale packet forwarding while maintaining low latency and high determinism. Their evaluation showed that VPP significantly outperforms traditional kernel-based stacks in throughput and scalability, especially under heavy traffic loads. These results establish VPP as a promising candidate for transport-layer optimization in network architectures where low jitter and processing efficiency are critical.

In their conference presentation at FOSDEM 2019, Kinsella and Konstantynowicz [13] demonstrated the use of FD.io CSIT (Continuous System Integration and Testing) to benchmark the performance of VPP in production-grade environments. Their results confirmed that VPP can reliably offload TCP and UDP packet processing from the kernel, achieving deterministic throughput even under varying load

conditions. The CSIT framework also enables automated, reproducible performance testing on high-speed hardware. However, despite these capabilities, the integration of CSIT-validated VPP architectures into fixed broadband infrastructures – such as HFC networks – remains largely unexplored.

Despite the considerable progress in upgrading HFC networks through DOCSIS 4.0, DAA, SDN/NFV, and user-space transport frameworks such as VPP and QUIC, the reviewed literature lacks experimental studies that integrate these technologies into a unified model tailored for HFC environments. In particular, there are no publicly validated results on deploying QUIC-over-UDP termination at the network edge using VPP within asymmetrical HFC infrastructures. This gap confirms the relevance of the proposed study, which aims to experimentally evaluate such a model under realistic traffic and hardware constraints using open benchmarking frameworks.

Materials and methods of research. The study is based on experimental modeling of QUIC protocol terminati-

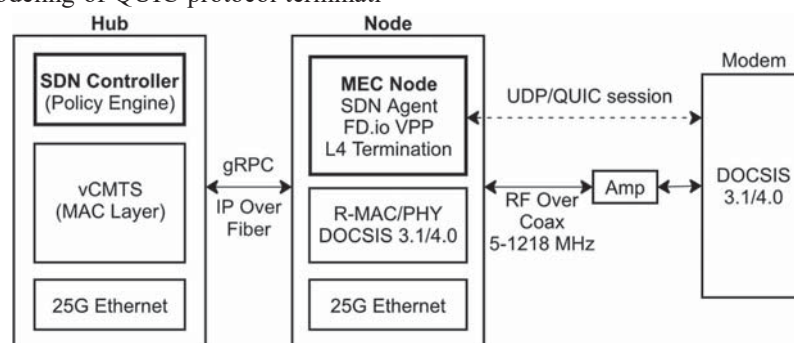


Figure. Edge-native HFC architecture with integrated QUIC termination / Edge-native архітектура HFC з інтегрованою термінацією QUIC

Figure illustrates the created architecture, which is logically decomposed into core and access domains. Each component serves a distinct role:

- **SDN Controller (Policy Engine):** Deployed at the operator's hub or datacenter, this component defines per-flow scheduling, congestion control, and queuing behavior. It communicates with the MEC node via a secure control plane interface (e.g., gRPC), dynamically adjusting traffic shaping policies through VPP's programmable pipeline.
- **vCMTS (MAC Layer):** Co-located with the SDN controller, the virtual Cable Modem Termination System retains all DOCSIS MAC-layer functionality, including service flow registration, bandwidth grants (MAPs), and configuration delivery. It interfaces with the RPD using standard Remote PHY protocols, while delegating all IP- and transport-layer processing to the MEC.
- **25G Ethernet Interfaces:** Both the core (hub) and access nodes are connected via 25G Ethernet links, which ensure sufficient bandwidth for DOCSIS 3.1/4.0 data streams and signaling. These links carry IP payloads and gRPC-based control messages between the SDN controller, vCMTS, and MEC.
- **MEC Node:** Positioned at the access edge, this node embeds the FD.io VPP HostStack, a high-performance user-space TCP/IP stack. Applications are redirected to HostStack via LD_PRELOAD, enabling session termination, traffic classification, and per-flow shaping with minimal latency and CPU overhead.
- **R-MAC/PHY Node:** Integrates both DOCSIS physical-layer functions (e.g., modulation/demodulation, symbol synchronization) and MAC-layer components (e.g., service flow scheduling, grant map generation) at the access edge. Unlike a Remote PHY Device (RPD), which delegates MAC operations to a centralized virtual CMTS, the R-MAC/PHY performs local DOCSIS MAC processing, enabling reduced

on at the transport layer within HFC networks. The methodology involves the use of the VPP software stack and a benchmarking tool to simulate asymmetric traffic loads with different numbers of clients and sessions. All test scenarios were executed under fixed network parameters in a controlled environment. The goal was to evaluate throughput dynamics under single-core constraints and identify scalability limitations of user-space transport processing.

Research results and their discussion / Результати дослідження та їх обговорення

Edge-Native HFC Network Architecture. To address the research objective of developing a scalable HFC architecture with edge-based transport-layer termination, we created a software-defined system design that enables relocation of QUIC-over-UDP processing to the access edge. This model eliminates the need for fiber replacement or DOCSIS modem upgrades, while improving upstream responsiveness and scalability.

backhaul latency, improved upstream responsiveness, and enhanced scalability under high subscriber concurrency. This architecture allows partial decoupling from centralized MAC logic while maintaining full compatibility with DOCSIS 3.1/4.0 modems.

- **Amplifier and Coaxial Segment:** This legacy section of the HFC network remains unchanged. Analog RF signals (5-1218 MHz) are amplified and delivered over coaxial cables, preserving backward compatibility with all existing installations.
- **DOCSIS 3.1/4.0 Cable Modem:** The modem remains unaware of the edge-based QUIC termination. It continues to function normally within DOCSIS specifications and connects to the RPD without modification.

This modular design separates PHY/MAC and IP/transport responsibilities, allowing per-session optimization at the edge.

CSIT Infrastructure. The performance of user-space QUIC-over-UDP termination was experimentally assessed using the FD.io CSIT infrastructure, specifically the 3n-icx testbed optimized for high-throughput benchmarking of VPP-based systems. The testbed comprises three bare-metal servers (Server-Type-F3), each equipped with dual-socket Intel® Xeon® Platinum 8360Y CPUs (Ice Lake-SP), 512 GB DDR4 memory, PCIe Gen4 NVMe storage, and dual 100G network interfaces (Intel E810-CQDA2 or NVIDIA ConnectX-6 Dx) operating over DPDK with VFIO drivers [3]. All nodes are directly interconnected via 100GE links, enabling latency-accurate measurement and NUMA-aware workload placement.

The software stack included VPP 24.02+ built with HostStack and VCL (Vector Communication Library), LD_PRELOAD redirection (libvcl_ldpreload.so) for unmo-

dified socket interception, and DPDK 24.07 with NUMA optimizations. All tests were executed on Ubuntu 24.04 LTS using Robot Framework integrated with custom Python libraries for automation and telemetry capture. Performance metrics – such as aggregate throughput, retransmissions, CPU utilization, and core affinity – were collected and stored as structured artifacts.

The CI/CD framework ensured reproducibility by resetting the software environment before each iteration. YAML-based configurations and public Jenkins triggers allow independent researchers to replicate the test conditions, scenarios, and results without requiring privileged access to the infrastructure. These experiments form the empirical basis for evaluating edge-based QUIC transport termination efficiency in HFC architectures using software-defined, user-space networking components.

Traffic Generation. Traffic was generated using `vpecho`, a native benchmarking utility included in VPP. This tool simulates lightweight transport-layer request – response behavior over UDP, reflecting the structure of real-world QUIC exchanges. Operating in a client-server model, the application establishes connections from a VPP client instance to a remote `vpecho` server. Each session transmits fixed-size UDP frames (1280 bytes) in a unidirectional stream from client to server, emulating upstream-dominant flows typical in residential DOCSIS deployments.

The `vpecho` client configuration defines [2]:

- number of logical clients (each with a unique source IP and UDP port);
- number of sessions per client;
- frame size and transmission interval.

Test parameters were passed via VPP CLI and YAML-based automation templates. During execution, the server echoes each received frame back to the client, completing a full round-trip per session. Key performance metrics – such as throughput and per-session counters – were collected from internal VPP telemetry.

Scenario Definitions. To investigate the effect of client and session scaling on QUIC-over-UDP performance in HFC networks, we defined four distinct traffic scenarios (Table 1). Each scenario was designed to isolate the performance impact of either horizontal (client-side) or vertical (session-side) scaling using VPP HostStack and the `vpecho` benchmarking tool:

- **Scenario S1 – Baseline:** this configuration uses a single client initiating one session, serving as the reference scenario for evaluating throughput under ideal, non-concurrent conditions. It establishes a performance baseline for comparison with all other configurations.
- **Scenario S2 – Vertical Scaling:** in this session-heavy setup, a single client initiates ten concurrent QUIC sessions toward the server. This tests the effect of session multiplexing on a single logical endpoint, helping assess session management overhead and context switching within VPP HostStack.
- **Scenario S3 – Horizontal Scaling:** this configuration increases the number of logical clients to ten, each initiating only one session. Each client has a distinct source IP address and UDP port, simulating distributed upstream flows. This isolates the impact of source diversity and horizontal distribution on CPU cache locality and per-client metadata management.
- **Scenario S4 – Full Load:** the most intensive scenario combines both dimensions of scaling: ten clients, each opening ten sessions, resulting in 100 simultaneous QUIC flows. This setup simulates highly concurrent workloads and aims to identify throughput bottlenecks and scheduling contention under edge resource constraints.

Table 1. Scenarios with QUIC client and session scaling / Сценарії з масштабуванням QUIC клієнтів і сесій

Scenario	Description	Client	Sessions per Client
S1	Baseline	1	1
S2	Vertical scaling (session-heavy)	1	10
S3	Horizontal scaling (client-heavy)	10	1
S4	Full load	10	10

All tests were conducted using a fixed frame size of 1280 bytes and a single VPP worker thread to emulate constrained processing typical of access-edge nodes. The direction of traffic was unidirectional (client to server) to reflect real-world HFC asymmetry, where upstream performance is often a limiting factor.

Research Results. This section presents the quantitative results obtained from the execution of the four traffic scenarios (S1-S4) defined in Table 1. The evaluation focuses on aggregate throughput (in Mbps) as the primary performance metric, accompanied by interpretation of how varying client and session counts affect performance. Table 2 summarizes the measured throughput per scenario, as extracted from CSIT telemetry logs.

Table 2. Average and normalized throughput across QUIC-over-UDP scenarios S1-S4 / Середня та нормалізована пропускна здатність у сценаріях QUIC через UDP (S1-S4)

Scenario	Avg Throughput, Mbps	Std. Dev., Mbps	Normalized Throughput, %
S1 (base)	882.8	~1.23	100
S2 (1c10s)	586.5	~1.78	66.4
S3 (10c1s)	625.3	~1.52	70.8
S4 (10c10s)	529.7	~2.04	60.0

Scenario analysis:

- *S1* (1 client, 1 session) serves as the baseline and achieves the highest absolute throughput (882.8 Mbps), likely due to minimal session tracking overhead and optimal cache locality within a single VPP worker thread.
- *S2* (1 client, 10 sessions) shows a significant performance drop (66.4 %), indicating that vertically scaled, session-heavy configurations introduce overhead at the VPP session layer – particularly under constrained threading.
- *S3* (10 clients, 1 session) maintains slightly better throughput (70.8 %) than *S2*, suggesting that distributing connections across source IPs improves pipeline parallelism despite added routing complexity.
- *S4* (10 clients × 10 sessions) yields the lowest performance (60.0 %), consistent with expected contention effects at both session dispatch and memory access levels.

These results demonstrate that even on a single VPP worker core, QUIC-over-UDP can achieve over 880 Mbps under optimal conditions – specifically, in the baseline scenario (1 client, 1 session) with minimal session tracking, no context switching across cores. In more complex multi-client or multi-session scenarios, the throughput dropped to 60-71 % of the baseline. As shown in Table 2, throughput declines non-linearly as concurrency increases, revealing diminishing returns with higher session density.

Discussion of research results. Recent studies have highlighted the performance limitations of QUIC over high-speed networks. They [11] conducted a systematic examination of QUIC's performance, revealing that the UDP+QUIC+HTTP/3 stack suffers a data rate reduction of up to 45.2 % compared to the TCP+TLS+HTTP/2 counterpart. This degradation is attributed to high receiver-side processing overhead, particularly excessive data packets and QUIC's user-space acknowledgments. In contrast, our

study demonstrates that by leveraging VPP's user-space processing capabilities, it is possible to mitigate such overheads, achieving higher throughput even under asymmetric traffic conditions.

In the study [9], the authors conducted a detailed performance comparison between production deployments of QUIC and TCP across Google, Facebook, and Cloudflare. Their benchmarking methodology revealed that QUIC's performance advantages are not guaranteed by protocol design alone, but are highly dependent on the choice and tuning of congestion control algorithms. In particular, they showed that QUIC clients require non-trivial configuration tuning and that implementation-level bugs or misconfigurations can negate the protocol's theoretical advantages. These findings support our observations that real-world performance in edge-native QUIC deployments is shaped as much by system-level factors and CPU affinity as by the protocol stack itself, especially under high concurrency and resource-constrained conditions.

In their work on interoperability testing, Seemann and Iyengar [8] introduced an automated framework for validating conformance and compatibility among QUIC implementations. They highlighted how divergent behavior in packet formats, flow control, and recovery logic leads to inconsistent performance and interoperability failures. Their findings underscore the need for structured test frameworks that evaluate QUIC behavior under different operational conditions. Our work complements this approach by applying reproducible benchmarking of QUIC-over-UDP transport under controlled resource constraints, specifically within HFC networks. While Seemann and Iyengar focus on cross-implementation correctness, our study demonstrates how internal scaling factors – such as session concurrency and client context spread – can introduce throughput variability even in single-stack, standards-compliant deployments. Together, these findings advocate for the co-evolution of interoperability testing and performance validation in real-world access scenarios.

The work [10] highlights how the 4th Gen Intel® Xeon® Scalable processors, combined with the Data Streaming Accelerator (DSA), significantly reduce memory copy overhead in both kernel and user-space stacks. In particular, the VPP Hoststack benefits from shared memory acceleration via DMA-based batch processing, which improves throughput in transport-layer tasks such as QUIC termination. Our experiments confirm these findings: the use of VPP HostStack with shared memory APIs and LD_PRELOAD in a single-core setup consistently achieved high upstream throughput, even under session scaling. This confirms that the DSA-enabled VPP architecture is well-suited for low-latency, CPU-efficient packet processing in HFC networks with asymmetric traffic conditions.

CableLabs discussed the convergence of 5G and HFC networks [14], emphasizing the role of multi-tenancy at the edge. Their 10G Lab initiative showcases the potential of deploying multiple network functions on shared infrastructure. Our study aligns with this vision, demonstrating that integrating QUIC termination at the edge using VPP can enhance performance without necessitating significant infrastructure changes.

So, based on the results of the work performed, it is possible to formulate the following scientific novelty and practical significance of the research results.

Scientific novelty of the obtained research results – for the first time, an edge-native transport-layer model for HFC networks has been developed and experimentally validated using user-space QUIC termination based on VPP, providing software-defined traffic control under asymmetric load conditions; the model has been further advanced by designing reproducible client/session scaling scenarios automated within the CI/CD-enabled CSIT infrastructure, which enabled the identification of throughput thresholds under increasing concurrency within a single-core CPU environment.

Practical significance of the research results – the results of the study can be applied to develop and implement a scalable and cost-effective modernization strategy for existing HFC networks by introducing a software-defined transport-layer model based on VPP and QUIC, which enables higher upstream throughput and reduced latency without physical infrastructure replacement, and can be further integrated into MEC platforms for flexible deployment in next-generation broadband access systems.

Conclusions / Висновки

Developed and evaluated an edge-native architecture for HFC networks with QUIC transport over UDP, which enables transport-layer functionality to be relocated to the access edge without requiring full physical infrastructure replacement. Based on the results of the research the following main conclusions can be drawn.

1. The highest throughput was observed under minimal concurrency (1 client, 1 session), confirming the sensitivity of VPP session handling to scale and memory locality.
2. Performance degraded non-linearly with increased session and client count, emphasizing the importance of horizontal scaling and isolation.
3. Compared to centralized DOCSIS 3.1 and DAA deployments with vCMTS, the created model offers reduced upstreaming network traffic delay, better session responsiveness, and improved effective HFC network bandwidth utilization.
4. The architecture requires no firmware updates, RF reconfiguration, or modem replacement, making it highly viable for constrained upgrade environments.
5. Limitations of the current study – such as unidirectional traffic, single-core testing, and absence of application-layer benchmarking – suggest future directions for integration with DOCSIS telemetry and multithreaded VPP scaling.
6. The results confirm that integrating software-defined transport logic at the HFC edge presents a cost-effective, modular enhancement path toward symmetrical broadband readiness – particularly relevant where DOCSIS 4.0 or FTTH is not yet economically justified.

Reproducibility and Acknowledgments. We gratefully acknowledge the LFN and the FD.io community for providing access to the test infrastructure and tooling. As the FD.io ecosystem is fully open source, all experiments conducted in this study are fully reproducible, allowing researchers to replicate or extend our results under identical conditions.

References

1. Barach, D., Linguaglossa, L., Marion, D., Pfister, P., Rossi, D., & Pontarelli, S. (2018). High-speed software data plane via vectorized packet processing. *IEEE Communications Magazine*, 56(11), 97–103. <https://doi.org/10.1109/MCOM.2018.1800069>
2. FD.io. (n.d.). *VPP HostStack External Echo Client/Server*. The Linux Foundation. URL: <https://wiki.fd.io/view/VPP/HostStack/ExternalEchoClientServer>

3. FD.io CSIT. (n.d.). *FD.io data center testbed specifications: 3-node IceLake (3 n-icx)*. The Linux Foundation. URL: https://csit.fd.io/cdocs/infrastructure/fdio_dc_testbed_specifications/#3-node-icelake-3n-icx
4. Iyengar, J., & Thomson, M. (2021). *QUIC: A UDP-based multiplexed and secure transport* (RFC 9000). RFC Editor. <https://doi.org/10.17487/RFC9000>
5. Kreutz, D., Ramos, F. M. V., Verissimo, P., Rothenberg, C. E., Azodolmolky, S., & Uhlig, S. (2015). Software-defined networking: A comprehensive survey. *Proceedings of the IEEE, 103*(1), 14–76. <https://doi.org/10.1109/JPROC.2014.2371999>
6. Liu, Y., Amin, R., & Stemmler, H. (2023). Digital adoption: Accelerating postpandemic, yet a widening divide. *World Bank, Digital Progress and Trends Report 2023*, 1–21. URL: <https://openknowledge.worldbank.org/server/api/core/bitstreams/b6125358-fb1c-47bb-a6e7-4896b2152904/content>
7. Salinger, J., & Sigman, S. (2021). *Lessons from operating tens of thousands of Remote PHY Devices*. SCTE & NCTA. URL: <https://www.nctatechnicalpapers.com/Paper/2021/2021-lessons-from-operating-tens-of-thousands-of-remote-phy-devices>
8. Seemann, M., & Iyengar, J. (2020). Automating QUIC interoperability testing. In *Proceedings of the Workshop on the Evolution, Performance, and Interoperability of QUIC (EPIQ'20)* (pp. 8–13). Association for Computing Machinery. <https://doi.org/10.1145/3405796.3405826>
9. Yu, A., & Benson, T. A. (2021). Dissecting performance of production QUIC. *Proceedings of the Web Conference 2021 (WWW'21)*, 12 pages. ACM. <https://doi.org/10.1145/3442381.3450103>
10. Yu, P., Liu, M., Wang, J., Qiu, P., DiGiglio, J., & Coras, F. (2023). *FD.io VPP – Accelerate the host stack with 4th Gen Intel® Xeon® Scalable Processor technology guide*. Intel Corporation. URL: <https://networkbuilders.intel.com/docs/networkbuilders/intel-avx-512-high-performance-ipsec-with-intel-xeon-scalable-processor-technology-guide-1683018859.pdf>
11. Zhang, Q., Deng, H., Bai, X., Chen, Y., Yang, Y., & Ren, K. (2024). Understanding the performance of QUIC over high-speed networks. *Proceedings of the ACM Web Conference 2024 (WWW'24)*, 4301–4311. <https://doi.org/10.1145/3589334.3645323>
12. Cable Television Laboratories. (2022, October). *DOCSIS® 4.0 Physical Layer Specification (CM-SP-PHYv4.0)*. URL: <https://www.cablelabs.com/specifications/CM-SP-PHYv4.0>
13. Kinsella, R., & Konstantynowicz, M. (2019, February). *Building production-grade networking software with FD.io CSIT* [Conference presentation]. FOSDEM 2019, Brussels, Belgium. <https://doi.org/10.5446/44177>
14. CableLabs. (2022, October 12). *10G: Enabling future-ready networks*. Cable Television Laboratories, Inc. URL: <https://www.cablelabs.com/insights/10g-enabling-future-ready-networks>
15. Charter Communications. (2022, December 13). *Investor meeting presentation*. URL: <https://ir.charter.com/events/event-details/charter-communications-investor-meeting>
16. Cable Television Laboratories. (2023, April). *Data-over-Cable Service Interface Specifications: Physical Layer Specification (CM-SP-PHYv3.1-I20-230419)*. URL: <https://www.cablelabs.com/specifications/CM-SP-PHYv3.1>

I. Д. Іванець¹, В. К. Овсяк^{1,2}

¹ Національний університет "Львівська політехніка", м. Львів, Україна
² Келецький технологічний університет, м. Кельце, Польща

EDGE-NATIVE АРХІТЕКТУРА ГІБРИДНОЇ ОПТИЧНО-КОАКСІАЛЬНОЇ МЕРЕЖІ З РЕАЛІЗАЦІЄЮ QUIC-ПРОТОКОЛУ ЧЕРЕЗ UDP

Запропоновано модель архітектури оптично-коаксіальної мережі з edge-native реалізацією транспортного рівня, що передбачає перенесення термінації QUIC-протоколу на вузли доступу. Встановлено, що централізоване опрацювання процесів транспортного рівня у традиційних мережах не забезпечує необхідної продуктивності для сучасних сервісів з асиметричним трафіком і підвищеними вимогами до затримки передачі даних. Оцінено переваги реалізації термінації QUIC (англ. *Quick UDP Internet Connections*) протоколу поверх UDP (англ. *User Datagram Protocol*) без фізичної заміни телекомунікаційної інфраструктури мережі доступу. Експеримент проведено на базі програмного стеку FD.io VPP, інструменту vppecho та інфраструктури тестування продуктивності CSIT з використанням стенду ICX-3N-100GE (три вузли на базі Intel Ice Lake із 100GE інтерфейсами Intel E810CQ). Усі вимірювання виконано за фіксованого розміру мережевого фрейму 1280 байт та обмеження в одне процесорне ядро. Сценарії тестування автоматизовано з використанням методології CI/CD (англ. *Continuous Integration/Continuous Deployment or Delivery*) та відкритого шаблону конфігурації тестового середовища, що забезпечує відтворюваність результатів. Проведено порівняння чотирьох конфігурацій QUIC протоколу: базової (882.8 Мбіт/с), з 10 клієнтами по 1 сесії (586 Мбіт/с), з 1 клієнтом на 10 сесій (586.5 Мбіт/с) та з 10 клієнтами на 10 сесій (562 Мбіт/с). Виявлено, що базовий режим роботи забезпечує найвищу пропускну здатність, тоді як масштабовані сценарії тестування демонструють очікуване зниження продуктивності роботи внаслідок зростання накладних витрат. Це зниження продуктивності зумовлене витратами на облік одночасно активних QUIC-сесій, впливом NUMA-архітектури та перемиканням контекстів між клієнтами в умовах обмеженого однопотокового середовища, що вказує на важливість стратегій масштабування транспортного рівня. Охарактеризовано потенціал створеної edge-native моделі як програмно-визначеної концепції для поетапної модернізації оптично-коаксіальних мереж доступу без потреби дороговартісної фізичної перебудови її. Реалізована модель сумісна з наявними мережами DOCSIS 3.1 і не потребує оновлення абонентського обладнання чи зміни кабельної інфраструктури. Її гнучкість і програмно-визначений характер забезпечують придатність до змінних вимог сучасних ширококосмугових сервісів із низькою затримкою трафіку та потребою в масштабованості навантаження.

Ключові слова: делегування транспортного рівня; симетричний ширококосмуговий доступ; високопродуктивне оброблення пакетів; масштабоване навантаження; програмно-визначені мережі; порівняльне тестування конфігурацій.