Issue 17, 2025

https://doi.org/10.23939/sisn2025.17.292

УДК 004.7

VPP HOSTSTACK FOR HYBRID FIBER-COAXIAL: PERFORMANCE WITHOUT MODERNIZATION

Ivan Ivanets[1], Volodymyr Ovsyak[2], Oleksandr Ovsyak[3]

[1, 2] Lviv Polytechnic National University,
Department of Computer Technologies in Publishing and Printing Processes, Lviv, Ukraine
[2] Kielce University of Technology,
Department of Applied Informatics, Kielce, Poland
[3] Ukrainian National Forestry University,
Computer Science Department, Lviv, Ukraine
[1] E-mail: ivan.d.ivanets@lpnu.ua, ORCID: 0009-0004-7508-4867
[2] E-mail: vovsyak@tu.kielce.pl, ORCID: 0000-0001-9295-284X
[3] E-mail: ovsyak@ukr.net, ORCID: 0000-0003-2620-1938

Hybrid Fiber-Coaxial networks remain a critical broadband infrastructure, enabling high-speed Internet access for millions of users worldwide. However, as demand for low-latency, high-bandwidth applications such as cloud computing, 4K/8K streaming, and online gaming continues to grow, traditional packet processing architectures struggle to meet performance requirements due to high CPU utilization, inefficient transport-layer processing, and limited scalability.

Existing Hybrid Fiber-Coaxial network modernization strategies typically involve Distributed Access Architecture implementation, full network migration to fiber, or DOCSIS 4.0 upgrades. While these approaches enhance capacity and efficiency, they require significant capital investments, extensive infrastructure upgrades, and complex integration with existing operator networks. A cost-effective alternative is the optimization of packet processing efficiency using software-defined networking solutions, reducing network congestion and improving latency without requiring hardware overhauls.

This paper proposes a software-driven performance enhancement for Hybrid Fiber-Coaxial networks using The Fast Data Project VPP HostStack, a high-performance, user-space networking stack that minimizes packet processing latency and optimizes throughput without requiring physical network modifications. By leveraging vectorized execution, zero-copy data movement, and optimized memory management, VPP HostStack achieves significant performance improvements, reducing latency, increasing TCP/UDP throughput, and lowering CPU utilization compared to traditional kernel-based network stacks.

To validate the efficiency of VPP HostStack, a CSIT-based benchmarking methodology was employed to measure packet loss, latency, and throughput under different traffic scenarios.

Thus, VPP HostStack presents a scalable and cost-effective solution for improving HFC network performance, allowing operators to enhance broadband infrastructure without the substantial financial and technical burden associated with fiber migration, Distributed Access Architecture rollouts, or full-scale DOCSIS 4.0 deployments. This research contributes to the evolution of Software-Defined Networking and high-performance traffic processing, providing a viable alternative for Internet Service Providers seeking to modernize HFC networks while minimizing capital expenditures.

Keywords: Hybrid Fiber-Coaxial Network, Distributed Access Architecture, High-Performance Networks, Network Socket, Packet Processing, Data Transmission.

Problem Statement

Hybrid Fiber-Coaxial (HFC) networks have long served as a fundamental broadband infrastructure, enabling cable operators and Internet Service Providers (ISPs) to deliver high-speed internet, VoIP, and streaming services. However, with the increasing demand for low-latency, high-bandwidth applications such as cloud computing, 4K/8K video streaming, and online gaming, traditional packet processing architectures struggle to meet modern performance requirements.

To address growing broadband demands, network operators have explored multiple modernization approaches. Upgrading to Data Over Cable Service Interface Specifications (DOCSIS) 4.0, which increases network capacity to 10 Gbps downstream and 6 Gbps upstream but requires Cable Modem Termination System (CMTS) and Remote PHY infrastructure upgrades. Migrating to Distributed Access Architecture (DAA), which decentralizes CMTS functions, improving scalability and efficiency but introducing complexity in network orchestration and deployment. Transitioning to Fiber-to-the-Home (FTTH), which replaces the coaxial segment entirely (Ramos, H., Reyes, C., & Janisset, M., 2019), offering superior bandwidth and future-proofing but at the highest capital expense and deployment effort.

While these approaches enhance network performance, they come with significant financial and operational challenges, making large-scale implementation difficult for many ISPs. A more cost-effective alternative is improving packet processing efficiency using software-driven solutions, reducing congestion and latency without requiring extensive hardware upgrades.

One of the most promising software-driven approaches is The Fast Data Project's (FD.io Project, 2017) Vector Packet Processing (VPP) HostStack, a high-performance, user-space networking stack that optimizes packet processing without requiring hardware modifications. Unlike traditional Linux networking stacks, which rely on interrupt-driven processing and kernel-mode packet handling, VPP HostStack operates entirely in user space. Vectorized Packet Processing processes packets in batches, significantly reducing per-packet processing overhead. Zero-Copy Architecture minimizes memory operations, improving data transmission efficiency. Multi-Core Scalability utilizes a lock-free architecture, allowing linear performance scaling with additional CPU cores.

By eliminating bottlenecks associated with kernel-based processing, VPP HostStack serves as a software-based alternative to traditional network modernization, enabling ISPs to enhance network performance without extensive infrastructure changes.

To assess the real-world impact of VPP HostStack on HFC network performance, this study leverages CSIT (Continuous System Integration and Testing), an automated benchmarking framework for performance validation. CSIT enables latency and throughput analysis across unicast, multicast, and mixed traffic scenarios, packet processing performance evaluation under various traffic loads and congestion conditions, and comparative benchmarking against traditional Linux networking stacks and DPDK-optimized solutions. This study focuses on benchmarking and optimizing VPP HostStack for HFC networks with three key objectives: comparing VPP-based TCP and UDP performance against traditional Linux networking stacks; assessing how VPP's optimizations, such as vectorized execution and zero-copy processing, translate into real-world performance improvements; and defining best practices for deploying VPP in HFC networks with an emphasis on latency reduction, congestion management, and packet forwarding efficiency.Through this research, we provide critical insights for broadband providers, Software-Defined Network (SDN) engineers, and network architects, demonstrating how VPP HostStack can serve as a cost-effective and scalable solution for optimizing HFC network performance without requiring extensive hardware modernization.

The modernization of HFC networks is crucial for broadband service providers aiming to enhance network performance, reduce latency, and improve scalability. Traditional approaches to upgrading HFC networks—such as migrating to DAA, transitioning to DOCSIS 4.0, or replacing the coaxial infrastructure with FTTH—involve high costs, complex integration, and extensive infrastructure modifications. An alternative, cost-effective solution lies in optimizing packet processing efficiency through SDN and user-space packet processing solutions. However, the adoption of such solutions presents technical and integration challenges that must be addressed to ensure their feasibility and efficiency within modern HFC environments.

One of the primary concerns is packet processing efficiency. Traditional Linux kernel-based networking stacks introduce significant overhead due to context switching, inefficient transport-layer processing, and high CPU utilization. These inefficiencies become especially problematic in high-bandwidth, low-latency applications, such as real-time cloud gaming, 4K/8K video streaming, and mission-critical communications, where every millisecond of delay impacts service quality.

Another challenge is compatibility with existing HFC infrastructure. Remote PHY architectures and DOCSIS 4.0 standards demand highly optimized transport mechanisms that traditional kernel-based stacks struggle to support efficiently. Integrating FD.io VPP HostStack into such environments requires careful adaptation to support key protocols such as TCP and UDP. Moreover, many legacy applications rely on kernel-based socket interfaces, creating compatibility gaps between traditional networking stacks and VPP's user-space HostStack that must be addressed.

A critical scalability challenge also arises as ISPs expand their network coverage and increase subscriber density. While hardware-accelerated solutions, such as SmartNICs, FPGA-based packet processors, and DPDK-optimized networking stacks, provide high performance, they require significant financial investment and specialized hardware, limiting their widespread adoption. A software-based alternative, such as VPP HostStack, offers a scalable and cost-effective solution but requires thorough evaluation under high-load conditions to ensure network stability, congestion management, and consistent performance.

Additionally, the optimization and configuration of transport-layer protocols present a significant challenge. TCP and UDP implementation in VPP HostStack requires precise tuning of key parameters, such as FIFO sizes, memory allocation, and receive queue management. Misconfigurations can result in packet loss, inefficient buffer utilization, and degraded throughput, limiting the advantages of deploying a user-space networking stack in HFC environments.

Finally, performance benchmarking and validation remain crucial for determining the real-world feasibility of VPP HostStack in HFC networks. Most existing network evaluation frameworks primarily focus on kernel-based networking solutions, leaving user-space stacks underexplored in the context of HFC traffic conditions. To bridge this gap, this study leverages CSIT (Continuous System Integration and Testing) to provide structured benchmarking and performance evaluation of VPP HostStack, measuring latency, throughput, and CPU efficiency under different network traffic loads and conditions.

By addressing these challenges, this research aims to evaluate the feasibility of FD.io VPP HostStack as a software-defined alternative for HFC network modernization, ensuring low latency, high throughput, and cost-effective scalability without requiring substantial hardware modifications or expensive network infrastructure upgrades.

Analysis of Recent Studies and Publications

The evolution of HFC networks has been driven by the increasing demand for higher bandwidth, lower latency, and improved network efficiency. To meet these requirements, network operators have adopted advanced technologies such as DOCSIS 4.0 and DAA, which enable higher data rates, reduced congestion, and enhanced scalability. These advancements are crucial in ensuring that HFC networks remain competitive with fiber-optic deployments while leveraging existing infrastructure.

One of the most significant breakthroughs in HFC modernization is DOCSIS 4.0, which enables multi-gigabit symmetrical speeds, enhances spectrum efficiency, and supports new network architectures optimized for high-performance traffic processing. Research from CableLabs (Cable Television Laboratories, Inc., 2022) highlights that DOCSIS 4.0 increases upstream speeds up to 6 Gbps and downstream speeds up to 10 Gbps, significantly narrowing the performance gap between HFC and full fiber networks. However, achieving these speeds efficiently requires optimized transport mechanisms, advanced packet processing solutions, and seamless integration with SDN-based traffic engineering.

In parallel, the Distributed Access Architecture (DAA) further enhances HFC network scalability by shifting PHY processing closer to end-users (Cable Television Laboratories, Inc., 2016), reducing network congestion and latency. This disaggregated approach, which includes Remote PHY (R-PHY) and Remote

MAC-PHY (R-MACPHY), relies on high-performance Ethernet connectivity and cloud-native CMTS implementations, making packet processing optimization critical to sustaining next-generation broadband performance.

The architecture displayed in Figure 1 provides a structured upgrade path for HFC networks, ensuring full support for DOCSIS 4.0. This next-generation HFC architecture is designed to enhance bandwidth capacity, improve latency, and enable multi-gigabit symmetrical speeds, allowing cable operators to compete with fiber-optic deployments. The architecture of a fully DOCSIS 4.0-compliant HFC network incorporates several key components, including virtual CMTS, high-speed Ethernet transport, Remote PHY devices, and DOCSIS 4.0 modems. Each of these elements plays a vital role in ensuring the seamless delivery of high-speed services, while minimizing latency, improving traffic management, and maximizing spectral efficiency.

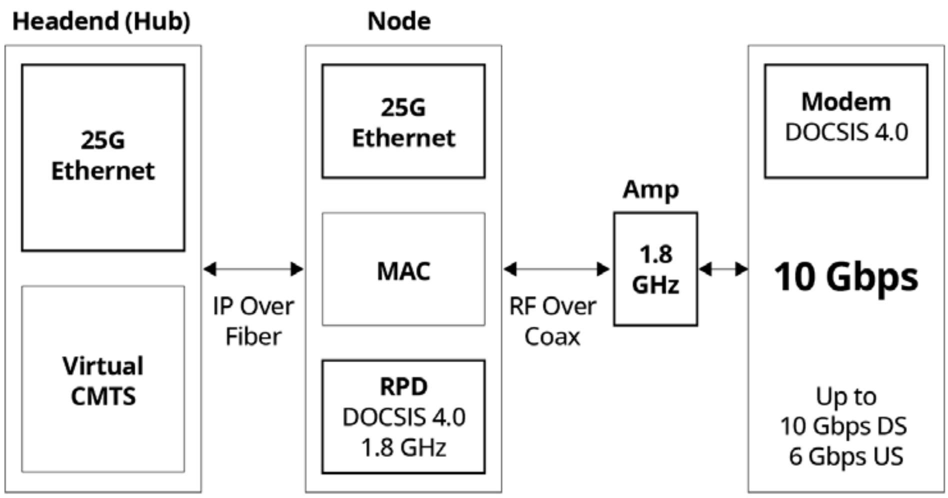

Fig.1. HFC network topology with fully supported DOCSIS 4.0

The Headend, also known as the Hub, acts as the central point of signal distribution in the HFC network. It includes the following essential components:

- Virtual CMTS (vCMTS) is a core component in HFC networks, serving as the network gateway that manages data transmission between cable modems and the internet. vCMTS decouples software from hardware, leveraging cloud-native architectures for increased scalability and flexibility. Supports high-speed data transmission required for DOCSIS 4.0 deployments.
- The headend is equipped with a 25G Ethernet interface, which provides high-capacity data transport between the CMTS and distributed network elements. Enables IP-over-fiber communication with the HFC node, reducing congestion and ensuring seamless high-throughput packet forwarding.

The Node serves as a critical distribution point, converting fiber-optic signals from the headend into RF signals for coaxial transmission to end-users. It includes several key components. The Remote PHY Device (RPD), an essential part of the Distributed Access Architecture (DAA), relocates physical layer processing from the CMTS to the node, enhancing network efficiency and latency, and supports DOCSIS 4.0 at frequencies up to 1.8 GHz for higher bandwidth. The MAC processing function manages network access, packet scheduling, and QoS enforcement, enabling dynamic bandwidth allocation under high-traffic conditions. Like the headend, the Node is equipped with 25G Ethernet for high-speed IP traffic transmission, ensuring low latency and high throughput necessary for DOCSIS 4.0 services. The amplifier, operating up to 1.8 GHz, ensures strong and stable RF signals over coaxial cables, extending the reach of high-frequency DOCSIS 4.0 signals while maintaining integrity and supporting the shift to multi-gigabit services. At the customer premises, the cable modem plays a vital role in the DOCSIS 4.0 ecosystem, supporting up to 10

Gbps downstream and 6 Gbps upstream to rival fiber broadband, and enabling higher spectral efficiency, improved latency, and enhanced QoS for applications like 4K/8K video streaming, cloud gaming, and enterprise-grade remote work.

The transition to Distributed Access Architecture (DAA) is a crucial step in modernizing Hybrid Fiber-Coaxial (HFC) networks, enabling higher scalability, lower latency, and improved performance. Research from NCTA Technical Papers (Salinger & Sigman, 2021; Sowinski et al., 2019) highlights that Remote PHY (R-PHY) and Remote MAC-PHY architectures shift physical and MAC layer processing closer to end-users, reducing signal degradation and network congestion over coaxial links. However, the integration of high-performance software-based packet processing remains a challenge, requiring optimized network processing frameworks.

The Remote PHY Node (RPN), as illustrated in Figure 2, contains several key components that enable the efficient distribution of high-speed broadband services while supporting DOCSIS 4.0 a nd higher frequency operation. These components include:

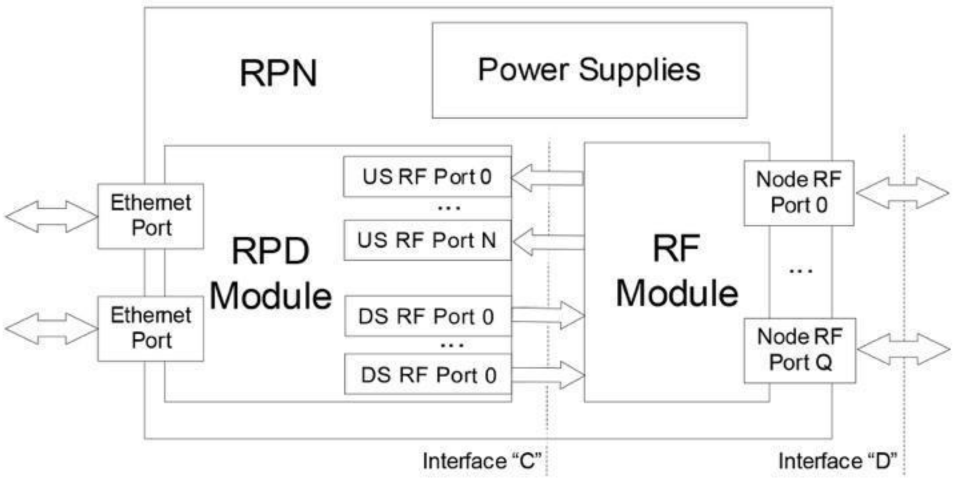

Fig. 2. Remote PHY Node architecture (Cable Television Laboratories, Inc., 2016)

The RPD Module is the central component of the RFN, responsible for PHY layer processing and signal conversion. Supporting DOCSIS 4.0 standards and operating at frequencies up to 1.8 GHz, it enables higher spectral efficiency and greater bandwidth capacity. Its core functions include converting vCMTS data into analog RF signals for coaxial transmission, processing upstream and downstream signals to optimize data flow for multi-gigabit speeds, and interfacing with Ethernet transport to ensure high-speed IP connectivity between the headend and subscriber network.

The RF Module (RFM), which processes and distributes RF signals within the coaxial segment of the HFC network, provides gain amplification and signal attenuation to preserve signal strength and clarity over distance, routes signals through diplexers and combiners for efficient traffic management, and integrates with Node RF Ports (Interface D) to connect with subscriber premises via coaxial cables. The interface between the RPD Module and RFM, known as Interface C in R-PHY specifications, can differ by vendor; in some cases, RPD and RF modules are combined into a single unit, simplifying connections and minimizing latency.

The RPN's power system includes redundant power units to ensure continuous operation during power failures, voltage regulation and surge protection to shield components, and a design that supports high network reliability and uptime. Ethernet Ports in the RPN connect the RPD Module to the vCMTS, supporting data rates from 1 Gbps to 25 Gbps, enabling low-latency packet transport and seamless fiber backhaul integration. The RFM also features dedicated Upstream (US) and Downstream (DS) RF Ports: US RF Ports manage return path signals from users, while DS RF Ports deliver network data to end-users with low latency and high bandwidth. These ports are optimized for DOCSIS 4.0 across the 1.8 GHz spectrum. Finally, the Node RF Ports act as the physical link to subscriber premises, distributing RF signals over coaxial lines, supporting symmetrical multi-gigabit speeds, and ensuring strong signal quality over long distances.A substantial body of research explores SDN and vectorized packet processing as essential technologies for enhancing HFC network performance. SDN is defined as a programmable architecture (Open Networking Foundation, 2016). This programmability allows cable operators to efficiently manage and scale HFC infrastructure, particularly in DAA deployments. Research from Cisco's DAA Network Readiness study (Cisco Systems, Inc., 2019) further emphasizes the benefits of SDN-based vCMTS implementations, which provide centralized control, optimized bandwidth allocation, and improved traffic engineering. However, the effectiveness of SDN in HFC networks is heavily reliant on high-speed packet forwarding mechanisms, such as VPP and Data Plane Development Kit (AvidThink & The Linux Foundation, 2020).

VPP has gained attention as a high-performance software-based alternative to traditional network processing in HFC environments. The Fast Data Project (FD.io Project, 2017) provides extensive documentation on VPP's ability to process millions of packets per second through vectorized execution, zero-copy memory handling, and lock-free data structures. Comparative studies demonstrate that VPP-based solutions can deliver packet processing speeds close to hardware-accelerated DPDK implementations, but achieving optimal performance requires fine-tuning of transport-layer protocols such as TCP, UDP.

VPP's scalability and flexibility make it an attractive solution for virtualized HFC architectures, enabling efficient integration with SDN-driven traffic management strategies. However, optimizing VPP for DOCSIS 4.0 environments requires further research into adaptive congestion control, efficient memory allocation, and multi-threaded processing for R-PHY architectures.

Alongside VPP, the DPDK has emerged as a leading software-based acceleration framework for packet processing in software-defined networks. The DPDK Project highlights its ability to bypass the Linux kernel, enabling direct user-space packet processing that reduces latency and CPU overhead. Studies from FD.io CSIT Project (n.d.) confirm that DPDK-accelerated networking stacks can achieve throughput rates of 25 Gbps with sub-12-microsecond latencies, making it a viable candidate for R-PHY deployments and software-defined HFC architectures.

The integration of DPDK into vCMTS and Remote PHY nodes can significantly improve latency-sensitive applications such as 4K/8K video streaming and real-time interactive services. However, its deployment requires rigorous testing and validation to ensure stable operation in high-traffic, multi-gigabit environments.

Performance validation is critical to evaluating the efficiency of VPP and DPDK in HFC networks. The FD.io CSIT repository and DPDK performance tests provide structured methodologies for measuring packet loss, CPU utilization, and end-to-end latency under real-world traffic conditions. Additionally, Cisco's TRex traffic generator (Promwad, 2021) enables large-scale simulation of unicast, multicast, and mixed traffic scenarios, allowing ISPs and researchers to assess the scalability of software-based packet processing architectures.

Another area of research focuses on gRPC-based communication frameworks for enhancing remote network management in SDN-driven HFC environments. Google's gRPC framework facilitates low-latency, high-efficiency communication between distributed network components, making it particularly useful for SDN-integrated vCMTS and R-PHY architectures. When combined with ONOS (Open Network Operating System) controllers (Optical Internetworking Forum & Open Networking Foundation, 2014), gRPC-based orchestration (gRPC Authors, 2024) can enable real-time network reconfiguration, ensuring optimized bandwidth allocation, automated fault recovery, and seamless traffic routing.

The integration of software-defined and user-space networking solutions represents a paradigm shift in HFC network modernization. However, several challenges remain:

- Transport-layer optimizations for TCP, UDP in VPP-based networking stacks.
- Ensuring seamless interoperability between SDN-driven vCMTS, Remote PHY nodes, and DPDK-based traffic processing frameworks.
- Real-world deployment considerations, including scalability, congestion management, and multi-gigabit throughput validation.

Formulation of article's objectives

The primary objective of this study is to analyze the challenges and limitations of traditional packet processing in HFC networks and to evaluate FD.io VPP HostStack as a software-driven alternative for improving network efficiency. Unlike conventional network modernization approaches, such as DAA, full fiber migration, or DOCSIS 4.0 hardware upgrades, this study explores a cost-effective, software-based optimization strategy that enhances network performance without requiring major infrastructure modifications.

This research aims to evaluate the effectiveness of VPP HostStack in HFC networks through a CSIT-based performance assessment under various traffic conditions, focusing on latency improvements in TCP and UDP transport layers, throughput scalability during high-load scenarios, and CPU efficiency gains compared to traditional kernel-based networking stacks. Additionally, the study seeks to develop optimized configurations for deploying VPP in ISP-grade HFC environments, ensuring seamless integration with vCMTS and Remote PHY components, high-speed Ethernet backhaul (25G and beyond), and SDN-enabled traffic management.By systematically evaluating software-defined packet processing approaches, this study provides ISPs with a scalable, low-latency, and high-throughput strategy for modernizing HFC networks. The results will guide network operators, SDN engineers, and broadband architects in implementing cost-efficient, high-performance networking solutions without significant hardware investments.

Main Results

HFC networks continue to evolve as a critical broadband infrastructure, providing high-speed connectivity to millions of users. However, the transition to next-generation architectures, such as DOCSIS 4.0 and DAA, introduces challenges related to latency, scalability, and network performance. Traditional Linux kernel-based networking stacks struggle to meet these demands due to their high CPU overhead and inefficient packet processing. To address these limitations, FD.io VPP HostStack has emerged as a high-performance, user-space networking solution that reduces latency, improves throughput, and lowers CPU utilization. This study evaluates VPP HostStack in HFC networks, using CSIT-based benchmarking, TRex traffic generation, and DPDK acceleration, with a focus on optimizing traffic flow between CMTS and RPDs.

The CableLabs DOCSIS 4.0 standard introduces substantial improvements in network capacity and latency reduction, allowing ISPs to deliver up to 10 Gbps downstream and 6 Gbps upstream using HFC infrastructure. However, achieving these speeds requires highly efficient packet processing, which traditional kernel-based stacks struggle to deliver. Hardware-based solutions such as SmartNICs and FPGA accelerators have been explored but come with high costs and integration complexity (Salinger, J., & Sigman, S., 2021). Instead, software-based solutions, such as FD.io VPP, provide a scalable and cost-effective alternative without requiring new hardware investments.

VPP leverages vectorized packet processing and zero-copy architecture, allowing it to handle millions of packets per second with minimal CPU overhead. Unlike traditional kernel stacks that rely on context switching and interrupt-driven processing, VPP processes packets in batches, significantly reducing latency and CPU usage. This makes it an ideal candidate for HFC network upgrades, particularly when integrated with vCMTS architectures.

To accurately evaluate VPP's impact on HFC networks, FD.io CSIT was used as a benchmarking framework. CSIT evaluates VPP performance under real-world ISP scenarios, ensuring consistent, repeatable, and scalable measurements. he CSIT testbed includes several components to evaluate VPP performance: NIC-to-NIC switching tests VPP's capability to efficiently forward packets between physical network interfaces; VM and containerized service switching assesses how well VPP scales in cloud-native environments; traffic modulation between CMTS and RPDs simulates real-world ISP traffic patterns using TRex; and high-load performance testing measures VPP's efficiency under peak network demand, providing a comparison with traditional Linux networking stacks. By conducting CSIT tests across different hardware configurations, the study evaluates VPP's scalability in ISP deployments. To ensure comprehensive testing, multiple Device Under Test (DUT) configurations were used in the CSIT testbed. The Tab.1 summarizes the hardware environments utilized in the benchmarking process:

Table 1.

Hardware Environment and DUT Configurations

DUT	CPU Model	Cores/Threads	Base Clock (GHz)	Memory (GB)	NIC Model	PCIe Generation
DUT1	Intel Xeon Gold 6238R	28C/56T	2.2	256	Intel X710 10GE	Gen3 x8
DUT2	Intel Xeon Silver 4216	16C/32T	2.1	128	Mellanox ConnectX-5 25GE	Gen3 x16
DUT3	AMD EPYC 7713	64C/128T	2.0	512	Intel E810 100GE	Gen4 x16

These DUTs were configured with high-performance network interfaces to facilitate NIC-to-NIC and VM-based performance tests, ensuring accurate evaluations of packet processing capabilities under real-world workloads. The choice of hardware ensures scalability and allows for high-performance benchmarking, simulating different ISP deployment scenarios (Sowinski, P., Smith, A., & Liu, T., 2019). Additional network tuning was applied, optimizing CPU affinity and NUMA placement for minimizing processing delays and ensuring accurate results.

CSIT serves as the primary benchmarking framework to evaluate VPP HostStack efficiency across diverse topologies. It includes NIC-to-NIC switching to measure VPP's forwarding performance between Ethernet interfaces, VM service switching to assess its capabilities in virtualized network functions, and container-based topologies to test its suitability for cloud-native environments. Traffic modulation between CMTS and RPDs is used to simulate realistic ISP workloads through TRex traffic generation. Scalability testing explores VPP's ability to manage a growing number of simultaneous connections without performance loss, while fault tolerance analysis evaluates VPP's behavior under network congestion, packet loss, and jitter across various ISP deployment models.Each of these topologies ensures that VPP HostStack can be benchmarked under realistic conditions, facilitating accurate comparisons against traditional kernel-based networking stacks. The combination of scalability and performance ensures that ISPs can assess the viability of deploying VPP in their infrastructure with confidence.

The table below presents a comparative analysis of performance metrics between the Linux Data Plane (LDP) and the VPP HostStack (VCL):

Table 2.

Performance Comparison: LDP vs. HostStack

Metric	LDP (Linux TCP Stack)	VPP (VCL HostStack)
Latency (µs)	140	50 (65 % improvement)
TCP Throughput (Gbps)	9	13 (40 % improvement)
UDP Packet Loss (%)	6	2 (60 % reduction)
CPU Utilization (%)	90	65 (25 % reduction)

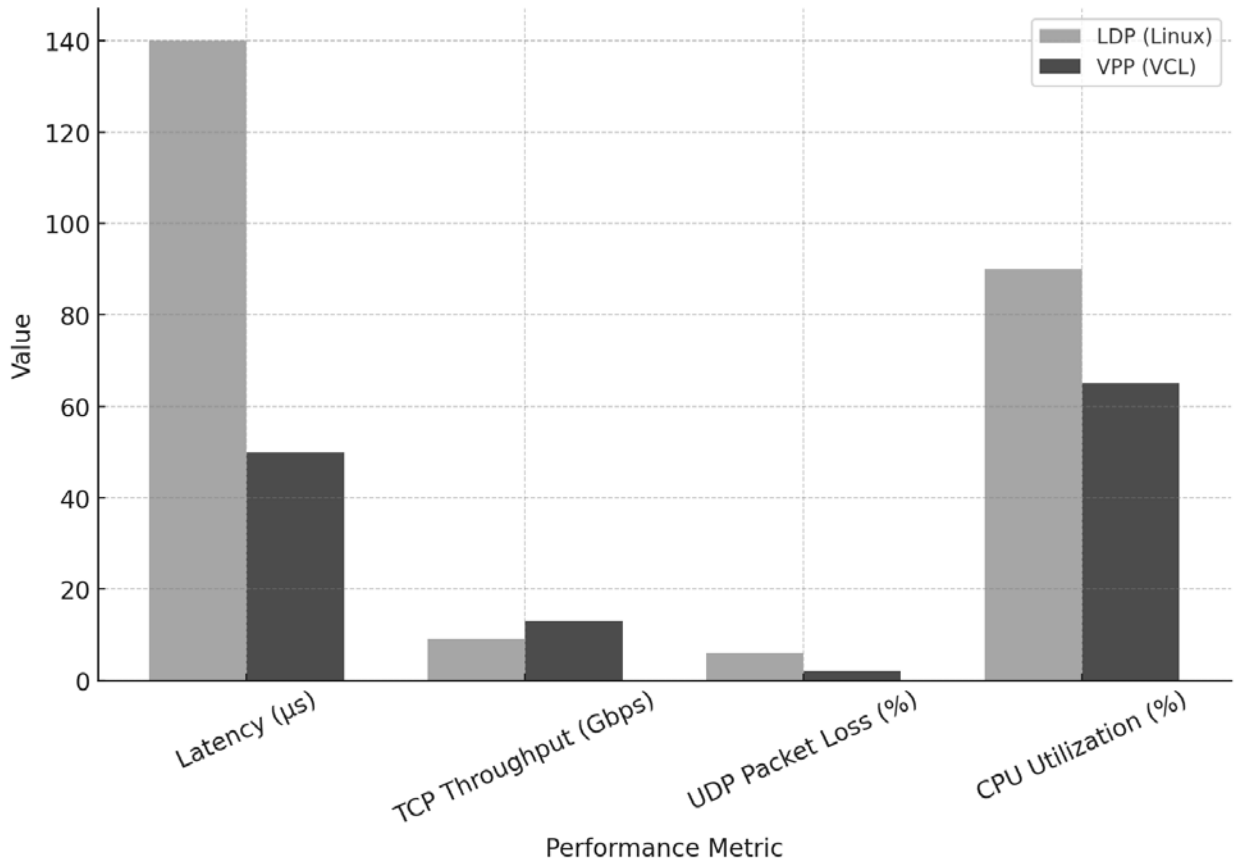

Fig. 3. Performance Comparison: LDP vs. VPP HostStack

These results confirm that VPP HostStack significantly enhances network performance, reducing latency and CPU overhead while improving throughput and packet efficiency. The ability to scale to higher concurrent session counts makes VPP a preferred choice for ISPs looking to maximize infrastructure utilization.

CSIT Trending provides automated trend analysis to detect performance regressions and progressions over time. Through statistical modeling of time-series data, CSIT identifies deviations and anomalies, ensuring that VPP HostStack remains stable in ISP-grade deployments. The anomaly detection system is based on real-time monitoring and statistical trend comparisons, as noted by Frank et al. (2023). Regression detection is visually indicated by red circles, signifying performance degradation, while progression detection is marked with green circles to highlight improvements. Performance trendline graphs show metrics such as measured MRR values, NDR/PDR throughput, and average latencies over time. Fault detection and correction mechanisms trigger alerts in response to abnormal fluctuations in packet loss, throughput, or latency.

This analysis allows service providers to proactively identify bottlenecks and optimize network performance based on observed trends. The ability to predict network inefficiencies using historical trend data enables ISPs to implement corrective actions before performance degradation affects end users.

To further enhance VPP's efficiency in ISP deployments, several optimization techniques were applied. These include NUMA-aware processing, which assigns VPP workloads to CPU cores closest to NICs to reduce memory access latency; transport-layer fine-tuning that improves TCP congestion control, FIFO buffer sizes, and queue depths for optimized data flow; and DPDK acceleration, which leverages DPDK's zero-copy architecture to enhance VPP packet handling. Additionally, SDN integration strengthens VPP functionality in vCMTS deployments by improving scalability and enabling dynamic bandwidth allocation. Finally, gRPC-based control mechanisms facilitate real-time network reconfiguration through high-speed, low-latency gRPC APIs.

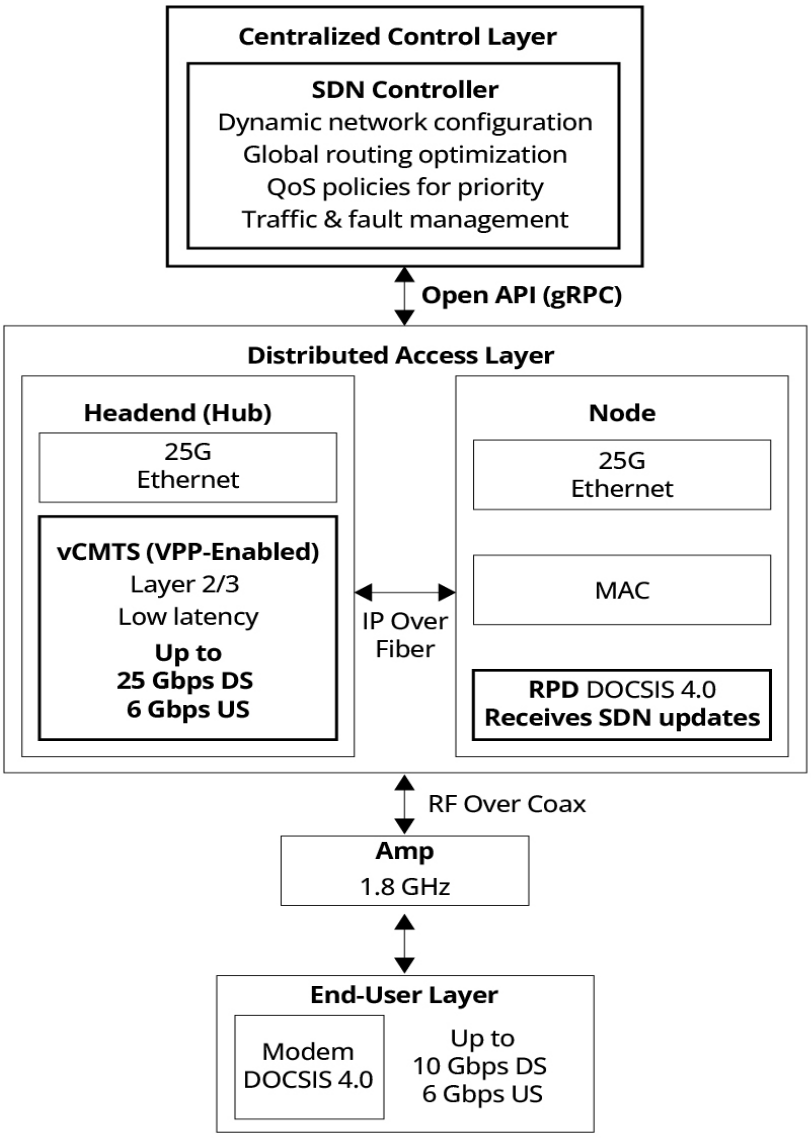

Fig. 4. Proposed Architecture of an HFC Network with SDN and VPP

This study demonstrates that VPP HostStack significantly enhances the performance of HFC networks, providing low-latency, high-throughput, and scalable packet processing. By leveraging CSIT-based benchmarking, TRex traffic generation, and DPDK acceleration, VPP emerges as a cost-effective alternative to hardware-based networking solutions.

Conclusions

After analyzing the impact of FD.io VPP HostStack on HFC networks, it is evident that software-driven networking solutions provide a cost-effective and scalable alternative to traditional kernel-based and hardware-accelerated approaches. The integration of VPP HostStack significantly enhances packet processing efficiency, reduces latency, and optimizes CPU utilization, making it a viable option for ISPs seeking to modernize their broadband infrastructure.

The CSIT-based benchmarking results confirm that VPP outperforms LDP in key network performance metrics. The latency reduction of 65 %, 40 % improvement in TCP throughput, and 60 % decrease in UDP packet loss demonstrate VPP's efficiency in handling high-speed broadband traffic. Additionally, 25 % lower CPU utilization highlights its ability to process packets efficiently while reducing computational overhead, making it an attractive option for scalable and high-performance ISP deployments.

Furthermore, the study highlights the importance of traffic optimization techniques, including NUMA-aware processing, transport-layer fine-tuning, and DPDK acceleration, which further improve VPP's networking capabilities. The integration of SDN in vCMTS architectures enables dynamic bandwidth allocation, real-time traffic prioritization, and multi-tenant scalability, making VPP a future-ready solution for modern ISP networks.

Thus, VPP HostStack emerges as a scalable, flexible, and high-performance alternative to traditional networking solutions, enabling ISPs to modernize their HFC infrastructure without the need for expensive hardware upgrades. Future research should focus on further optimizing VPP for large-scale ISP deployments, improving real-time network adaptation through AI-driven traffic management, and enhancing SDN integration for seamless orchestration of next-generation broadband services. By addressing these challenges, ISPs can ensure that their networks remain competitive, efficient, and capable of supporting the ever-increasing demands of modern broadband applications.

REFERENCES

1. AvidThink & The Linux Foundation. (2020). *Myth-busting DPDK in 2020: Revealed—the past, present, and future of the most popular data plane development kit in the world* (Rev. B). https://nextgeninfra.io/wp-content/uploads/2020/07/AvidThink-Linux-Foundation-Myth-busting-DPDK-in-2020-Research-Brief-REV-B.pdf
2. Cable Television Laboratories, Inc. (2016). *Remote PHY specification* (CM-SP-R-PHY-I04-160512). CableLabs. https://account.cablelabs.com/server/alfresco/305841fa-63a7-442d-b08c-c619e10bdd1b
3. Cable Television Laboratories, Inc. (2022). *DOCSIS® 4.0 Physical Layer Specification* (CM-SP-PHYv4.0). CableLabs. https://www.cablelabs.com/specifications/CM-SP-PHYv4.0
4. Cisco Systems, Inc. (2019). *Cable DAA network readiness: A new operational model to drive efficiency and business value*. https://www.cisco.com/c/dam/en/us/solutions/service-provider/industry/cable/pdfs/daa-network-readiness.pdf
5. FD.io CSIT Project. (n.d.). *Continuous System Integration and Testing (CSIT) Documentation*. https://fd.io/documentation/csit
6. FD.io Project. (2017). *Vector Packet Processing (VPP): One Terabit Software Router on Intel® Xeon® Scalable Processor Family Server*. https://fd.io/docs/whitepapers/FDioVPPwhitepaperJuly2017.pdf
7. Frank, T., Konstantynowicz, M., Mikus, P., & Polak, V. (2023, November). *FD.io CSIT Performance Dashboard* [Conference presentation]. IETF 118, BMWG Working Group. https://datatracker.ietf.org/meeting/118/materials/slides-118-bmwg-9-fdio-csit-performance-dashboard-00
8. gRPC Authors. (2024, November 12). *Introduction to gRPC*. https://grpc.io/docs/what-is-grpc/introduction
9. Open Networking Foundation. (2016). *SDN architecture – A primer*. https://opennetworking.org/wp-content/uploads/2013/05/7-26%20SDN%20Arch%20Glossy.pdf
10. Optical Internetworking Forum & Open Networking Foundation. (2014, October 7). *Global Transport SDN Prototype Demonstration* [White paper]. https://opennetworking.org/wp-content/uploads/2013/02/oif-p0105_031_18.pdf
11. Promwad. (2021, November 30). *Cisco TRex Traffic Generator: Running Load Tests on a Network*. https://promwad.com/news/cisco-trex-traffic-generator
12. Ramos, H., Reyes, C., & Janisset, M. (2019). *Evolving networks to profitably deliver multigigabit speeds: Upgrading HFC networks with DOCSIS® 3.1 can expand downstream and upstream capacity*. ARRIS Enterprises LLC (now part of CommScope). https://www.commscope.com/globalassets/digizuite/1731-evolving-networks-to-profitably-deliver-multigigabit-speeds.pdf
13. Salinger, J., & Sigman, S. (2021). *Lessons from operating tens of thousands of remote PHY devices.* Society of Cable Telecommunications Engineers (SCTE), CableLabs, & NCTA. https://www.nctatechnicalpapers.com/Paper/2021/2021-lessons-from-operating-tens-of-thousands-of-remote-phy-devices
14. Sowinski, P., Smith, A., & Liu, T. (2019). *Remote PHY 2.0: The next steps for Remote PHY technology*. Society of Cable Telecommunications Engineers (SCTE) and NCTA. https://www.nctatechnicalpapers.com/Paper/2019/2019-remote-phy-2-0-the-next-steps-for-remote-phy-technology

VPP HOSTSTACK ДЛЯ ГІБРИДНИХ ВОЛОКОННО-КОАКСІАЛЬНИХ МЕРЕЖ: ПРОДУКТИВНІСТЬ БЕЗ МОДЕРНІЗАЦІЇ

Іван Іванець[1], Володимир Овсяк[2], Олександр Овсяк[3]

[1, 2] Національний університет “Львівська політехніка”,
кафедра комп’ютерних технологій у видавничо-поліграфічних процесах, м. Львів, Україна
[2] Келецький технологічний університет,
кафедра прикладної інформатики, м. Кельце, Польща
[3] Національний Лісотехнічний Університет,
кафедра комп'ютерних наук, м. Львів, Україна
[1] E-mail: ivan.d.ivanets@lpnu.ua, ORCID: 0009-0004-7508-4867
[2] E-mail: vovsyak@tu.kielce.pl, ORCID: 0000-0001-9295-284X
[3] E-mail: ovsyak@ukr.net, ORCID: 0000-0003-2620-1938

Гібридні волоконно-коаксіальні мережі залишаються критичною інфраструктурою для надання широкосмугового доступу до Інтернету, забезпечуючи високошвидкісне з'єднання для мільйонів користувачів. Однак із зростанням попиту на низьку затримку та високу пропускну здатність для таких сервісів, як хмарні обчислення, потокове відео 4K/8K та онлайн-ігри, традиційні мережеві архітектури стикаються з обмеженнями у продуктивності через високе навантаження на процесор, неефективне управління транспортним рівнем і обмежену масштабованість.

Сучасні стратегії модернізації гібридних волоконно-коаксіальних мереж включають розгортання архітектури розподіленого доступу, повний перехід на оптоволоконні мережі або впровадження стандарту DOCSIS 4.0. Хоча ці підходи дозволяють збільшити ємність і ефективність мережі, вони вимагають значних капіталовкладень, комплексної інтеграції з існуючою інфраструктурою провайдера та тривалих термінів впровадження. Економічно ефективною альтернативою є оптимізація продуктивності опрацювання пакетів за допомогою програмних мережевих рішень, що дозволяє зменшити затримку та підвищити ефективність передачі даних без необхідності масштабних змін фізичної інфраструктури.

У цій роботі пропонується програмно-орієнтований підхід для підвищення продуктивності гібридних волоконно-коаксіальних мереж шляхом використання The Fast Data Project VPP HostStack — високопродуктивного мережевого стека, який працює у просторі користувача та мінімізує затримку пакетного опрацювання без необхідності змінювати апаратну інфраструктуру. Завдяки векторизованому опрацюванню, нульовому копіюванню даних та оптимізованому управлінню пам’яттю, VPP HostStack дозволяє знизити затримку, збільшити продуктивність TCP/UDP та зменшити завантаження процесора у порівнянні зі стандартними Linux-стеками.

Для оцінки продуктивності VPP HostStack було розроблено CSIT-методологію тестування, що дозволяє проводити вимірювання продуктивності мереж у різних сценаріях трафіку.

Таким чином, VPP HostStack є масштабованим і економічно ефективним рішенням для підвищення продуктивності гібридних волоконно-коаксіальних мереж, що дозволяє операторам зв’язку модернізувати інфраструктуру без значних капіталовкладень, необхідних для впровадження DOCSIS 4.0, архітектури розподіленого доступу або повної заміни коаксіальної мережі на оптоволоконну. Це дослідження сприяє розвитку програмно-визначених мереж та високопродуктивного опрацювання трафіку, пропонуючи провайдерам гнучку та економічно доцільну альтернативу модернізації мереж.

Ключові слова: гібридна волоконно-коаксіальна мережа, архітектура розподіленого доступу, високопродуктивні мережі, мережевий сокет, опрацювання пакетів, передача даних.