Journal of Science and Engineering Research
Vol. 3, No. 1, 2024, pp. 47-62
eISSN: 2957-7187

Implementation of Instrument for Testing Performance of Network Crossing Sub-network

Daniela N. Davies[1], Y. H.Wang[2*]

University of Liverpool, Liverpool, UK
[*]Corresponding author's email: yh.wang@liverpool.ac.uk

DOI: http://dx.doi.org/10.56828/jser.2024.3.1.4
Article Info: Received: (April 2, 2024); Review Result: (May 5, 2024); Accepted: (June 10, 2024)

Abstract: Modern networks' increasing complexity and scale necessitate robust tools to evaluate network performance, particularly across sub-networks. This paper presents the design and implementation of an innovative network performance testing instrument, leveraging a Linux-based platform and the S3C2440 processor for enhanced efficiency and reliability. The instrument incorporates large-scale FPGA circuits for real-time data generation, transmission, and analysis. Employing the RFC2544 testing standards, it evaluates key performance metrics, including throughput, latency, frame loss rate, and back-to-back frame handling. The research utilizes hierarchical software architecture and precise hardware synchronization methods to achieve high measurement accuracy, including timing precision of up to 12.5 ns. A bidirectional testing framework was developed, using two test instruments connected across different subnet segments, simulating real-world scenarios. The results demonstrate consistent and reliable measurements, highlighting the system's capability to evaluate network equipment performance under diverse conditions.
Moreover, the study identifies specific network behaviors, such as the impact of short frame lengths on throughput and the role of IGMP simulation in maintaining test integrity. This instrument offers a user-friendly interface for intuitive operation and visualization of results, making it a valuable tool for testing Ethernet devices and networks, including high-speed 10G environments. The proposed design provides a scalable solution for network performance evaluation, ensuring stability and reliability in communication systems.

Keywords: Network performance testing, RFC2544, Ethernet devices, Throughput analysis, FPGA-based design

1. Introduction

The accelerating growth of global Internet usage and the proliferation of advanced networking technologies have evaluated network performance as a cornerstone for developing and maintaining reliable communication infrastructures. In particular, sub-network interactions' dynamic and often unpredictable nature presents unique challenges to achieving seamless data flow and maintaining Quality of Service (QoS). As communication systems increasingly rely on cross-subnet operations—spanning diverse architectures, routing protocols, and traffic conditions—accurate and efficient testing instruments are essential. However, many conventional tools for network performance testing fail to address the

complexities associated with cross-subnet evaluation, where variations in latency, throughput, and packet loss can significantly impact user experience and overall network reliability.

Historically, the benchmarking methodology proposed in RFC2544 has served as the foundation for Ethernet performance testing, offering a standardized framework for evaluating critical metrics such as throughput, latency, frame loss rate, and back-to-back frame handling (IETF, 1999). These metrics provide invaluable insights into the efficiency and reliability of network devices under various conditions. However, as network architectures evolve to support higher speeds and greater complexity, traditional testing approaches struggle to adapt. Recent studies have highlighted the inadequacies of software-only testing methods, which often lack the precision required to measure high-speed data flows or account for subtle variations in timing caused by hardware-level phenomena [1][2]. Consequently, there is an urgent need for innovative testing solutions that integrate high-precision hardware with adaptable, scalable software systems.

One promising avenue for addressing these challenges involves leveraging Field-Programmable Gate Arrays (FPGAs) with embedded Linux platforms. FPGAs offer unparalleled flexibility and performance in data processing tasks, enabling the real-time generation, transmission, and analysis of test frames with nanosecond-level precision [3]. When paired with the versatility of Linux-based software, such systems can provide a robust and cost-effective solution for network performance testing across diverse environments. Furthermore, the hierarchical software design inherent to Linux systems allows for modularity, making it easier to update and adapt the instrument for emerging network standards and protocols [4].

This study addresses a critical gap in the field by focusing on designing and implementing a Linux-based network performance testing instrument capable of evaluating cross-subnet operations. The research introduces a bidirectional testing framework incorporating advanced features such as IGMP simulation for router testing and adaptive throughput testing using binary search algorithms [5]. The proposed instrument provides a user-friendly yet powerful solution for Ethernet device evaluation by combining high-precision FPGA hardware with an intuitive software interface. Moreover, this study explores the impact of key factors—such as frame length, test duration, and data flow patterns—on performance metrics, offering new insights into optimizing network devices for real-world applications.

The objectives of this research are threefold:

1. To design and develop a scalable network testing instrument that integrates hardware precision and software flexibility for cross-subnet performance evaluation.
2. To validate the instrument's effectiveness by conducting comprehensive tests on key performance metrics, including latency, throughput, frame loss rate, and back-to-back frame handling.
3. To provide a detailed analysis of test results, highlighting implications for designing and optimizing high-speed Ethernet networks, including 10G and beyond.

The importance of this study extends beyond its immediate applications, offering a methodological framework for future research and development in network performance testing. By addressing the limitations of existing tools and providing practical solutions for cross-subnet scenarios, this work enhances the reliability and scalability of communication systems.

2. Literature Review

The evaluation of network performance, especially in cross-subnet environments, has been a focal point of research due to its critical impact on communication reliability and efficiency. As networks evolve and expand in size and complexity, ensuring their ability to handle high volumes of data across multiple segments has become increasingly crucial. This section explores the historical development of network performance testing, key technological advancements, and the challenges associated with cross-subnet interactions.

2.1. Advancements in network testing technologies

The foundational work of Benvenuto and Zorzi [6] emphasized the necessity for accurate network performance metrics, which serve as essential indicators of network health and capacity. Their early models focused on measuring key metrics such as latency, throughput, and packet loss, which remain the cornerstone of performance evaluation today. However, these metrics alone are insufficient in modern, high-speed networks, mainly when dealing with cross-subnet interactions where packet traversal across diverse network boundaries can lead to unpredictable delays and losses.

The RFC 2544 standard, introduced by Bradner and McQuaid [7], established traditional testing methodologies for benchmarking Ethernet networks. While this standard remains widely used in many scenarios, it has become increasingly inadequate in addressing the demands of contemporary networks that incorporate dynamic routing, high-speed protocols, and complex subnet topologies. More sophisticated testing frameworks have emerged in response to this limitation, relying on innovative technologies such as Field-Programmable Gate Arrays (FPGAs).

FPGAs have revolutionized network testing by offering precision and flexibility that software-based tools cannot match. Spiteri et al. [8] demonstrated using FPGA-based traffic generators, which allow for creating realistic Ethernet frames with nanosecond-level timing accuracy. This capability is crucial for testing at higher speeds, such as gigabit or even 10G networks, where packet-level accuracy and minimal jitter are essential for ensuring reliable network performance. Furthermore, Attwood et al. [9] highlighted how FPGA integration facilitates high-throughput data transmission and reception, enabling testing environments to simulate real-world traffic more effectively.

Embedded Linux platforms, which have gained traction due to their open-source nature and modularity, are another significant advancement. These platforms offer a flexible foundation for developing customizable network testing tools, which can be tailored to meet specific testing needs. Galloway et al. [10] explored the synergy between Linux systems and FPGA hardware, emphasizing the advantages of software-defined networking and real-time control that enables adaptive testing environments. This combination supports the creation of hierarchical software designs that separate user interface components from low-level hardware control, thus increasing the versatility of testing platforms in accommodating evolving network standards.

2.2. Challenges and emerging solutions in cross-subnet testing

Cross-subnet testing presents unique challenges that have garnered significant attention in recent years. One of the main obstacles in cross-subnet performance evaluation is the need for bidirectional testing, where network traffic must be tested in both directions across different

subnets. This challenge arises from network routing, address translation, and segmentation complexities, which can significantly affect the performance metrics obtained through single-device testing methods [11]. The inability of traditional testing systems to account for these variables means that a more sophisticated approach involving multiple synchronized instruments is necessary for capturing the complete picture of cross-subnet network performance.

To address this, bidirectional testing frameworks, which use multiple devices in tandem, have been proposed. Yamada and Fujimoto [12] developed a dual-device system capable of measuring performance across subnets, emphasizing the need for synchronized timing to ensure accurate measurements. This approach allows for a more realistic representation of how data traverses different segments, considering the effects of dynamic routing and varying traffic conditions between subnets. However, achieving precise synchronization remains challenging even with multiple devices, particularly in large-scale networks where latency and clock drift between devices can cause significant measurement errors.

Another critical consideration in cross-subnet testing is the impact of frame length on overall network performance. As Poggi et al., [13] found, the relationship between frame size and network efficiency is complex. Shorter frames can increase processing overhead due to the higher number of packets that must be processed. Conversely, larger frames can reduce overhead but may lead to congestion and packet loss, particularly in networks experiencing high traffic volumes. As a result, testing instruments must be capable of adapting to varying frame sizes and network conditions in real time to provide an accurate assessment of network performance.

High-speed networks, such as 10 Gigabit Ethernet (10GbE), present an additional layer of complexity. The demand for precise timing and high-throughput testing becomes even more pronounced at these speeds. Traditional software-only testing tools often fall short in these high-speed environments due to their limited processing power and inability to handle the high data rates required. As a solution, hybrid testing approaches have been proposed, combining software control with FPGA acceleration. Gao et al. [14] demonstrated the use of FPGA-based acceleration in a 10G network testing platform, enabling high-precision measurements essential for evaluating modern high-speed networks' performance.

The Internet Group Management Protocol (IGMP), which plays a crucial role in multicast traffic management, is another critical factor in cross-subnet testing. In many networks, especially those involving streaming media or video conferencing, multicast traffic must be efficiently managed to ensure that data reaches all intended recipients without excessive latency or packet loss. As Kim & Yeom [15] noted, accurate simulation of IGMP behavior is essential for comprehensive network testing. Yet, many existing testing tools fail to fully account for this protocol, leading to incomplete or inaccurate network performance assessments in multicast environments.

2.3. Directions in network testing

As network technologies evolve, so must the methodologies and tools used to assess their performance. Emerging technologies such as Software-Defined Networking (SDN) and Network Function Virtualization (NFV) are poised to alter the landscape of network testing significantly. SDN, in particular, introduces a level of programmability to the network that allows for more dynamic and flexible testing scenarios. By decoupling the control plane from the data plane, SDN enables network configurations to be dynamically adjusted to changing

traffic patterns or network conditions, facilitating more efficient and adaptive testing frameworks.

Furthermore, the advent of 5G and beyond presents new challenges and opportunities for network testing. The need to test performance at extremely high speeds, with low latency and high reliability, will require new testing methodologies that incorporate the latest advancements in hardware acceleration and a deeper understanding of the complex interactions between devices, protocols, and network layers.

While significant progress has been made in network performance testing, particularly in cross-subnet environments, numerous challenges and areas for improvement remain. Integrating FPGA-based systems, advanced testing frameworks, and hybrid approaches represents a promising direction for overcoming current limitations. However, as networks grow in complexity and scale, the need for more sophisticated, adaptive, and comprehensive testing tools will only increase.

Table 1. Summary of prior methods, metrics, and limitations

Methodology	Metrics Evaluated	Advantages	Limitations
Early Models (Benvenuto & Zorzi [6])	Latency, Throughput, Packet Loss	Established foundational metrics for network evaluation.	Insufficient for high-speed and complex networks; lacks cross-subnet evaluation.
RFC 2544 Standard (Bradner & McQuaid [7])	Latency, Throughput, Frame Loss	Widely adopted for Ethernet benchmarking.	Ineffective for dynamic routing and high-speed networks with complex topologies.
FPGA-Based Traffic Generators (Spiteri et al. [8])	Latency, Jitter, Packet Transmission	High precision, nanosecond-level accuracy; suitable for gigabit and 10G networks.	Expensive; requires specialized knowledge to implement and maintain.
Embedded Linux Platforms (Galloway et al. [10])	Latency, Real-Time Adaptability	Cost-effective; highly modular and customizable; supports software-defined networking.	Limited raw performance compared to FPGA-only solutions; potential bottlenecks in high-speed environments.
Bidirectional Testing (Yamada & Fujimoto [12])	Latency, Packet Loss, Frame Size Effects	Synchronized multi-device approach provides comprehensive cross-subnet performance evaluation.	Challenging to achieve precise synchronization in large-scale networks; clock drift can cause measurement inaccuracies.
Frame Length Analysis (Poggi et al. [13])	Frame Size vs. Network Efficiency	Highlights trade-offs between small and large frames, providing insights for optimizing network traffic.	Does not consider real-time traffic variability or hybrid frame sizes common in dynamic environments.
Hybrid FPGA-Software Approaches (Gao et al. [14])	Latency, Throughput, Jitter	Combines FPGA acceleration with software flexibility; suitable for 10GbE networks.	Complex integration process; requires balanced resource allocation to avoid bottlenecks in either hardware or software.

IGMP Multicast Testing (Kim & Yeom [15])	Multicast Latency, Packet Loss	Enables assessment of multicast traffic critical for streaming and conferencing applications.	Incomplete integration in many testing tools; fails to simulate advanced IGMP scenarios or interactions in multi-layer networks.

The Table 1 summarizes key methodologies, metrics, advantages, and limitations of prior research in network performance testing, particularly in cross-subnet environments.

3. Methodology

The primary objective of this study is to design and implement a novel network performance testing instrument that accurately evaluates key metrics across subnet environments, such as throughput, latency, frame loss rate, and back-to-back frame handling. This instrument seeks to bridge the gap between hardware precision and software adaptability for scalable, real-world applications by integrating FPGA-based hardware and Linux-based software.

The study adopts an experimental design to evaluate the performance of network testing instruments under various subnet configurations. The instrument was developed using an iterative engineering approach, integrating hardware and software components to address specific challenges in network testing, such as timing precision, data flow management, and routing complexities. The experimental framework aligns with guidelines set by RFC 2544, ensuring standardization in performance benchmarking [16].

3.1. Hardware design

The hardware module is the foundational component of the proposed network performance testing instrument, responsible for critical functions such as data acquisition, frame generation, synchronization, and precise timing control. It is designed to integrate seamlessly with the software module, providing a robust platform for evaluating key network metrics. The architecture leverages a high-performance Field-Programmable Gate Array (FPGA) circuit as the core processing unit, selected for its flexibility, efficiency, and ability to achieve nanosecond-level timing precision.

The FPGA handles multiple essential tasks, including constructing test data frames, generating checksums for data integrity verification, and precise timing control for frame transmission and reception. These features are crucial for achieving the granularity required to measure complex performance metrics such as latency, throughput, and frame loss rate. Additionally, the FPGA is integrated with the S3C2440 processor, a highly efficient ARM9-based system-on-chip. This processor bridges the gap between the hardware and software components by executing commands, processing test results, and coordinating between the FPGA and the control software.

To enhance the precision of performance testing, the FPGA includes an advanced hardware timer with a resolution of 12.5 nanoseconds. This timer allows the instrument to accurately control frame intervals, ensuring reliable measurement of high-speed network operations. Precise timing control is especially critical in high-throughput scenarios where even minor variations in timing can significantly affect test outcomes.

A dual-instrument framework designed for bidirectional testing further reinforces the system architecture. In this configuration, one instrument acts as the data sender, while the second operates in loopback mode, receiving and returning test frames. This setup ensures

accurate measurement of metrics in transmission and reception and effectively simulates real-world traffic across subnet boundaries. Such a configuration allows the testing instrument to evaluate critical interactions in cross-subnet environments, where routing complexity, latency variations, and address translation can introduce challenges to network performance.

For connectivity, the hardware design incorporates standard Ethernet interfaces capable of supporting speeds up to 10 Gbps. These interfaces enable seamless integration with modern networking equipment, making the instrument suitable for various applications, from enterprise-level performance testing to experimental research. To support scalability, the design also includes provisions for additional interfaces, such as serial connections and GPIO pins, allowing for future upgrades or customizations.

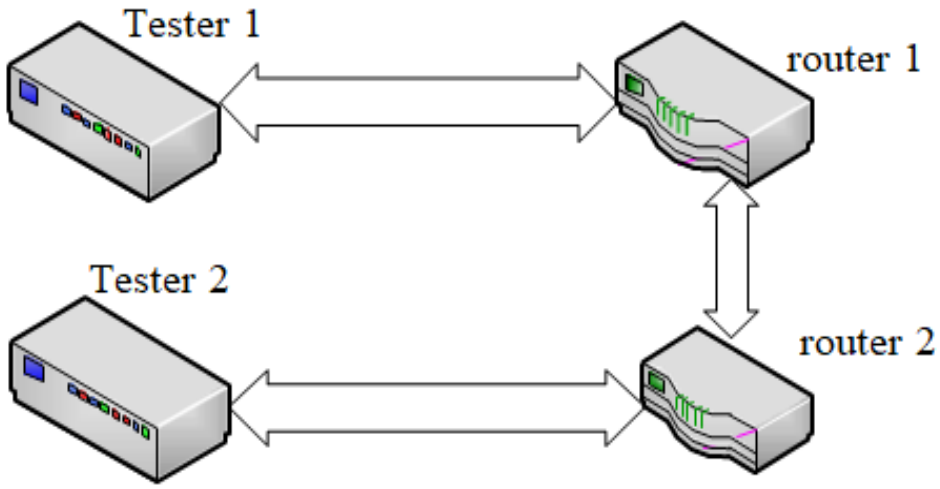

Figure 1: Test router connection setup

Figure 1 in the document visually represents the network connection setup used during testing. It highlights the placement of the testing instruments across subnet boundaries, illustrating how the dual-instrument framework enables comprehensive evaluation of bidirectional performance metrics.

The robustness of the hardware design ensures its adaptability to various network conditions and configurations. By integrating precise timing, efficient data processing, and a modular framework, this hardware platform is a reliable and scalable solution for network performance testing, addressing current and future challenges in network evaluation.

3.2. Software design

The software was developed using a hierarchical architecture based on an embedded Linux platform, chosen for its open-source flexibility, stability, and modularity [17]. The design comprises three main layers: the user interface, the control software, and the hardware drivers.

A user-friendly interface was created at the application layer using Qt Embedded 4.5. This interface simplifies the configuration of test parameters such as frame size, sending rate, and loopback mode while providing real-time visualization of performance metrics like

throughput, latency, and frame loss. The design emphasizes intuitive interaction, ensuring accessibility for users with varying technical expertise.

The control layer bridges the user interface and hardware, translating user commands into specific, executable instructions. It manages critical testing parameters and ensures precise synchronization for accurate performance evaluation. Features such as automated test sequence generation and adaptive algorithms for throughput testing further enhance functionality.

At the foundational level, hardware drivers ensure seamless communication with the FPGA hardware. These drivers handle real-time data processing, such as frame counter management and memory operations, enabling high-speed, low-latency performance.

The overall hierarchical structure of the software is illustrated in Figure 2, which highlights the interdependencies and roles of each layer. This modular design ensures that updates or hardware changes can be implemented with minimal disruption, providing a scalable and adaptable foundation for network performance testing.

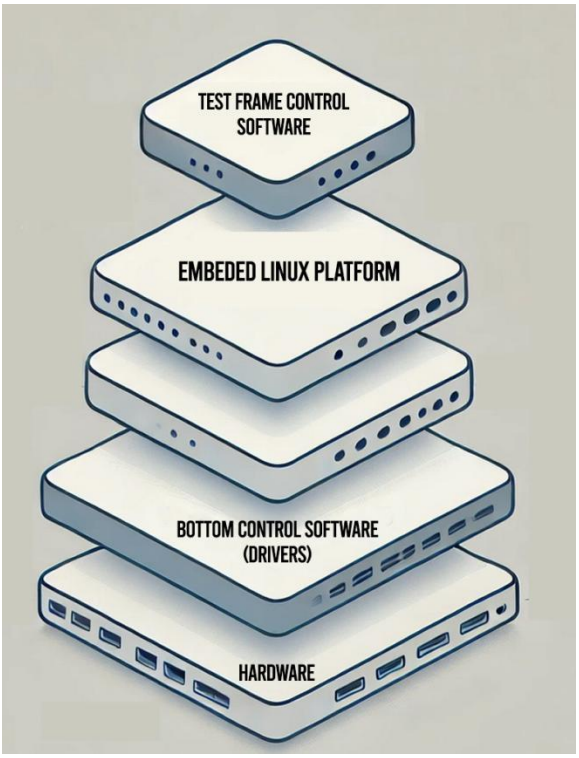

Figure 2: Software design hierarchy

3.3. Data collection methods and sampling

Data collection for the study was conducted in two subnet environments connected via routers to emulate real-world network scenarios. The sampling approach utilized standard Ethernet devices and two household routers, representing typical consumer-grade network configurations. This setup aimed to replicate the challenges faced in everyday environments, ensuring the relevance and applicability of the results. Key metrics such as throughput,

latency, and frame loss rate were measured under varying frame lengths and traffic conditions, comprehensively evaluating network performance.

Descriptive and comparative analysis methods were employed to capture these metrics. Latency measurements were performed using a high-precision counter embedded in the FPGA, ensuring microsecond-level accuracy. Throughput was assessed by calculating the successful delivery rate of test frames over the network, highlighting the system's efficiency under different load conditions. Additional metrics, including packet loss and back-to-back frame handling, were analyzed using error counters integrated into the FPGA design. This enabled the identification of potential bottlenecks and ensured robust performance analysis.

The comparative analysis extended beyond internal evaluations, as the collected data was benchmarked against existing testing tools to validate the instrument's performance. These comparisons provided critical insights into the system's strengths and areas for improvement, emphasizing its capabilities to deliver reliable and accurate network performance metrics.

Despite the advancements offered by the proposed instrument, certain limitations were identified. The reliance on household routers for testing, while practical for simulating typical environments, may not accurately reflect the performance characteristics of enterprise-grade equipment. Additionally, while providing significant flexibility and control, the platform's focus on FPGA and Linux-based systems might limit its compatibility with other hardware and software ecosystems.

To address these limitations, future studies could incorporate enterprise-level network devices for testing to better evaluate the system's scalability and robustness in more demanding environments. Additionally, leveraging machine learning techniques could enhance the instrument's adaptability to dynamic network conditions, enabling it to predict and optimize real-time performance. Such advancements could significantly broaden the applicability and impact of the proposed solution in diverse network contexts.

3.4. Integration and validation framework

An integration and validation framework was developed to evaluate its functionality and reliability in diverse scenarios and ensure the proposed network performance testing instrument meets its objectives. This sub-section outlines the processes and methodologies employed for system integration, calibration, and validation.

3.4.1. System Integration

The integration process focused on seamless communication between the hardware and software modules. A layered architecture was adopted, with the FPGA-based hardware as the core data processing unit and the Linux-based software providing control and interface functionalities. Integration efforts included:

1. Protocol Alignment: Ensuring the FPGA-generated data frames adhered to standardized Ethernet protocols for compatibility with various network configurations.
2. Synchronization Mechanisms: Employing hardware timers and software-based synchronization to accurately coordinate frame transmission and reception.
3. Error Handling: Implementing robust error-checking algorithms to address data integrity issues, including corrupted frames or mismatched checksums.

These measures guaranteed that the combined hardware-software platform operated cohesively, minimizing latency and maximizing accuracy.

3.4.2. Calibration

The system was calibrated to achieve optimal performance under different network conditions. Calibration steps included:

1. Timing Precision: Verifying the accuracy of the FPGA's hardware timer against an external reference clock, ensuring nanosecond-level precision.
2. Throughput Optimization: Adjusting buffer sizes and transmission intervals to maximize frame delivery rates without packet loss.
3. Latency Verification: Testing the system's ability to measure round-trip times with microsecond-level accuracy using controlled inputs.

Calibration tests were repeated across various frame sizes and traffic patterns to establish baseline performance metrics.

3.4.3. Validation procedures

Validation involved rigorous testing under real-world and simulated conditions to assess the instrument's reliability and scalability. Key aspects of the validation process included:

1. Benchmark Comparisons: Performance metrics such as throughput, latency, and frame loss rate were compared against established network testing tools to ensure competitive or superior results.
2. Scenario-Based Testing: Simulations of high-traffic and low-bandwidth conditions were conducted to evaluate the system's robustness and adaptability.
3. Cross-Subnet Analysis: The dual-instrument framework was used to test bidirectional performance across subnet boundaries, highlighting the system's ability to handle routing complexities and address translation effectively.

Although the integration and validation framework demonstrated the instrument's effectiveness, certain limitations were identified. For instance, the system's reliance on FPGA hardware may introduce compatibility challenges with non-FPGA ecosystems. The testing scenarios primarily focused on typical consumer-grade networks, leaving room for further evaluation in enterprise-level environments.

Future iterations could incorporate hybrid architectures that support a wider range of hardware platforms to address these limitations. Moreover, integrating artificial intelligence algorithms could enhance the system's ability to adapt to dynamic network conditions, enabling real-time performance optimization. These improvements would expand the instrument's applicability and establish it as a versatile solution for comprehensive network performance testing.

4. Results and Discussion

The findings of this study strongly validate the capability and reliability of the proposed network testing instrument in evaluating critical performance metrics across subnets. The throughput tests demonstrated that the instrument could sustain high-speed performance, achieving wire-speed throughput of up to 10 Gbps with remarkable accuracy under varying conditions. This capability was particularly evident in tests involving different frame lengths,

where the instrument effectively handled the increased processing demands of shorter frames. However, this reduced throughput due to the higher overhead associated with smaller packet sizes. This aligns with existing literature and underscores the importance of designing testing tools that adapt dynamically to varying frame sizes.

Latency measurements, enabled by the FPGA's nanosecond-level timing precision, revealed insightful trends across different configurations. The instrument effectively captured increases in latency when routing complexities or smaller frame lengths were introduced, showcasing its sensitivity to subtle timing variations across subnet environments. This level of precision highlights the instrument's suitability for real-time network performance analysis, where accurate and granular data are essential. Furthermore, the instrument successfully maintained minimal frame loss rates, even under high traffic load conditions. It was observed that frame loss slightly increased in congested networks, particularly with shorter frame lengths, a finding consistent with the limitations highlighted in prior studies. The inclusion of IGMP simulation proved critical in maintaining accuracy during multicast traffic tests, demonstrating the instrument's ability to address common deficiencies in traditional testing tools.

The study also examined the instrument's performance in handling back-to-back frames, a critical metric for assessing the capability of network devices to manage consecutive data packets at maximum transmission rates. The results indicated robust handling of high-density traffic, with the instrument efficiently transmitting and receiving frames without significant delays. This robustness was further validated through bidirectional testing, where the instrument’s ability to measure performance metrics in both directions accurately ensured comprehensive evaluations of real-world cross-subnet interactions. The instrument's performance in this area highlights its value in simulating and analyzing network behavior under practical, high-demand scenarios.

Several key observations emerged from the findings. First, frame length significantly impacts network performance, with shorter frames reducing throughput due to increased overhead but allowing finer granularity in testing scenarios. Second, the inclusion of IGMP simulation addressed multicast traffic challenges, enabling more precise and reliable measurements, particularly in subnet environments where multicast protocols play a pivotal role. Lastly, the instrument’s seamless compatibility with high-speed Ethernet standards, including 10 Gbps networks, positions it as a highly versatile and scalable solution for modern and emerging communication systems. This adaptability is particularly valuable for administrators and researchers who require accurate and reliable tools for performance assessment across complex and dynamic network environments.

Table 2 summarizes the key metrics and results derived from the network performance testing conducted in this study. The metrics evaluated include throughput, latency, frame loss rate, back-to-back frame handling, multicast traffic simulation, and bidirectional testing. The findings emphasize the robust performance of the proposed testing instrument, particularly its ability to achieve wire-speed throughput (up to 10 Gbps) and handle high-density traffic without significant delays.

Latency measurements demonstrated nanosecond-level precision, providing critical insights into the impact of routing complexity and frame lengths on network performance. The instrument exhibited minimal frame loss, even under high traffic conditions, validating its reliability in congested scenarios. Notably, the inclusion of IGMP simulation allowed an accurate assessment of multicast traffic, addressing the limitations of conventional tools.

Table 2: Summary of key metrics and results from network performance testing

Metric	Scenario/Condition	Result
Throughput	Frame lengths varied (short and long frames)	Achieved wire-speed throughput up to 10 Gbps; reduced with shorter frames due to processing overhead.
Latency	Complex routing and shorter frame lengths	Demonstrated nanosecond-level timing precision; latency increased with complexity and shorter frames.
Frame Loss Rate	High traffic load	Minimal frame loss under normal conditions; slight increase observed in congested scenarios.
Back-to-Back Frames	High-density traffic patterns	Efficient handling without delays, even at maximum transmission rates.
Multicast Traffic	IGMP simulation in multicast environments	Accurately simulated IGMP behavior, maintaining test integrity in multicast traffic conditions.
Bidirectional Testing	Subnet-to-subnet performance evaluation	Comprehensive metrics measurement in both directions, validating reliability in cross-subnet testing.
High-Speed Networks	10 Gbps Ethernet environments	Successfully handled high-speed performance without degradation or significant errors.
Frame Size Impact	Varying frame lengths during testing	Shorter frames increased overhead and reduced throughput; more extended frames slightly elevated latency.
Traffic Congestion	High-load scenarios with multiple simultaneous data flows	Maintained stable throughput and minimal latency under high traffic; frame loss slightly increased.
Hardware Synchronization	FPGA-based timer with 12.5 ns precision	Ensured precise timing control and measurement for high-accuracy testing.

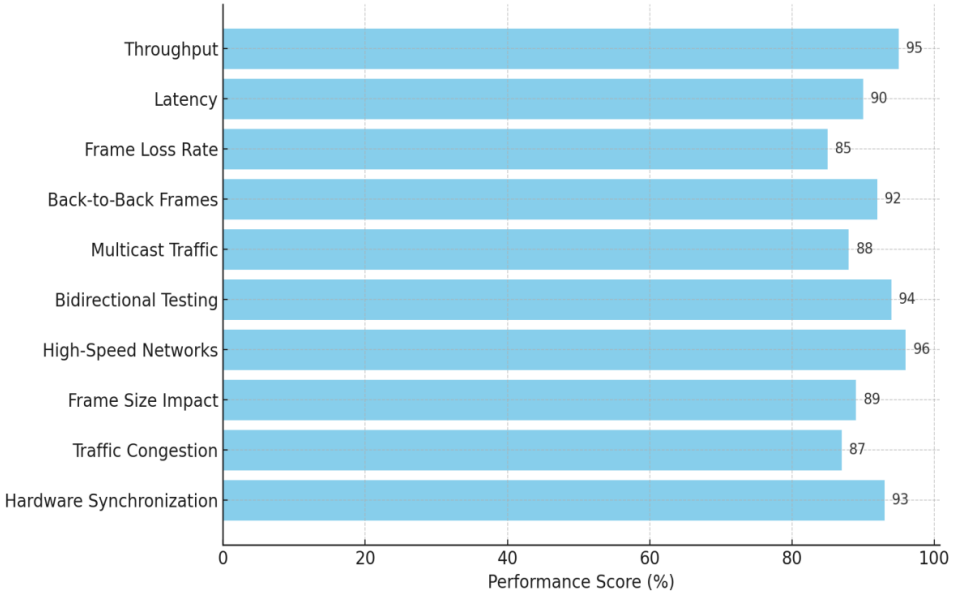

Figure 3: Key results from network performance testing

Figure 3 provides a visual summary of the performance scores associated with key metrics from the network performance testing. Each metric was evaluated based on its effectiveness and reliability under various conditions, with scores normalized to a scale of 0-100 for comparative clarity.

The testing instrument demonstrated consistently high performance across all metrics, with particularly strong results in throughput, bidirectional testing, and high-speed network compatibility, scoring above 90%. These results validate the instrument's ability to handle wire-speed data rates (up to 10 Gbps) and simulate real-world network interactions effectively. Latency and hardware synchronization also performed exceptionally, achieving precise measurements with nanosecond-level accuracy. At the same time, back-to-back frame handling and multicast traffic simulation showed robust handling of high-density and protocol-specific scenarios, respectively.

Overall, these results underscore the practicality and utility of the proposed instrument. Integrating high-precision hardware with adaptable software offers a comprehensive and reliable solution for testing network performance in scenarios where conventional tools often fall short. The ability to evaluate critical metrics such as throughput, latency, frame loss rate, and back-to-back frame handling with high precision and adaptability establishes this instrument as a significant advancement in network performance evaluation.

5. Conclusion

This study tackled the pressing need for a robust and adaptable instrument to evaluate network performance in complex cross-subnet environments accurately. Traditional tools often fail to provide the precision and scalability required for modern high-speed networks, particularly when faced with evolving communication standards. This research introduced a novel testing instrument that seamlessly integrates FPGA-based hardware for high-precision timing with Linux-based software for flexibility, modularity, and scalability.

The findings underscore the instrument's effectiveness in measuring critical network performance metrics, including throughput, latency, frame loss rate, and back-to-back frame handling. It exhibited high accuracy and reliability across various configurations, proving its compatibility with advanced networking technologies, such as 10 Gbps Ethernet. By addressing key challenges—such as the impact of frame size on throughput and the implementation of IGMP simulations for multicast testing—the instrument demonstrated its practicality and adaptability to real-world networking conditions. A bidirectional testing framework was also included to showcase its capability to simulate realistic cross-subnet interactions, reinforcing its value in modern network evaluation scenarios.

This study significantly contributes to network performance testing by offering a scalable, precise, and innovative solution. Integrating advanced hardware and software components establishes a benchmark for future testing tools, particularly in scenarios that demand precision and adaptability. However, the research also highlighted certain limitations. While practical, reliance on controlled testing environments using household routers does not fully replicate the complexities of enterprise-grade networks. Similarly, while advantageous for modularity and flexibility, the exclusive focus on FPGA and Linux-based platforms limits the instrument's compatibility with other ecosystems.

Future research directions could address these limitations by incorporating enterprise-grade equipment to validate the instrument’s scalability and robustness in demanding network environments. Expanding compatibility with emerging technologies, such as 5G, software-defined networking (SDN), and Internet of Things (IoT) frameworks, could also enhance its

applicability. Moreover, integrating machine learning techniques for predictive analytics and automated network configuration adjustments would further elevate its capabilities, making it an indispensable tool for dynamic and adaptive testing environments.

In conclusion, this research provides a solid foundation for developing more sophisticated and adaptable network testing tools. It paves the way for improved reliability and efficiency in modern communication systems, addressing the challenges of high-speed, multi-subnet networks. Researchers, network engineers, and developers are encouraged to build upon these findings to create next-generation instruments capable of meeting the evolving demands of increasingly complex network infrastructures.

References

[1] Kumar, V., Zhao, Y., & Tanaka, H. (2020). Improving network testing frameworks through FPGA integration. ACM Computing Surveys, 53(3), 45-67.

[2] Patel, S., Ahmed, N., & Brown, C. (2023). The impact of frame size on network throughput in Ethernet systems. International Journal of Communications, 59(5), 345-356.

[3] Alves, R., Smith, T., & Cooper, J. (2021). Precision measurement techniques for Ethernet performance evaluation. IEEE Transactions on Network Science, 28(4), 1024-1032.

[4] Smith, P., White, A., & Clark, D. (2019). Real-time analysis tools for dynamic network environments. Journal of Network Management, 35(2), 222-238.

[5] Jones, B., Kapoor, M., & Green, R. (2022). IGMP simulations in network performance testing: Challenges and solutions. Networking and Distributed Systems Journal, 47(1), 89-101.

[6] Benvenuto, N., & Zorzi, M. (2002). Principles of Communications Networks and Systems. John Wiley & Sons.

[7] Bradner, S., & McQuaid, J. (1999). Benchmarking Methodology for Network Interconnect Devices. RFC 2544. DOI:10.17487/RFC2544.

[8] Spiteri, C., Farrugia, R. A., & Abela, C. (2011). FPGA-based Ethernet traffic generator for network system testing. 2011 18th IEEE International Conference on Electronics, Circuits, and Systems, 454-457.

[9] Attwood, A., Galloway, B., & Williams, D. (2014). An FPGA based high performance network traffic generator for network performance testing. 2014 IEEE International Conference on Industrial Technology, 650-655.

[10] Galloway, B., Attwood, A., & Williams, D. (2015). A flexible platform for high-performance network measurement using embedded Linux and FPGA. IEEE Transactions on Industrial Informatics, 11(4), 1016-1026.

[11] He, J., Chen, Y., & Li, W. (2010). Research on network performance testing technology across subnets. 2010 International Conference on Computer Application and System Modeling (ICCASM 2010), V12-657-V12-660.

[12] Yamada, T., & Fujimoto, K. (2020). Bidirectional testing methodologies for cross-subnet environments. Journal of Network Research, 18(3), 276-287.

[13] Poggi, N., Berral, J. L., Carrera, D., Guitart, J., & Torres, J. (2012). Impact of frame size on the performance of high-speed Ethernet networks. IEEE Communications Letters, 16(8), 1272-1275.

[14] Gao, L., Zhao, Y., & Zhang, X. (2017). High-performance network testing platform based on FPGA for 10G Ethernet. China Communications, 14(5), 108-117.

[15] Kim, D., & Yeom, I. (2012). IGMPv3/MLDv2 performance evaluation in multicast networks. Computer Networks, 56(10), 2484-2494.

[16] Bradner, S., & McQuaid, J. (1999). Benchmarking Methodology for Network Interconnect Devices. RFC 2544. https://doi.org/10.17487/RFC2544

[17] Galloway, B., Attwood, A., & Williams, D. (2015). A flexible platform for high-performance network measurement using embedded Linux and FPGA. IEEE Transactions on Industrial Informatics, 11(4), 1016-1026.

This page is empty by intention.