Experimental Validation of On-Time Forwarding in Long-Haul Optical Networks for 6G Services

Changho Choi
Network Research Division
Electronics and Telecommunications Research Institute (ETRI)
Daejeon, South Korea
chhchoi@etri.re.kr

Yeoncheol Ryoo
Network Research Division
Electronics and Telecommunications Research Institute (ETRI)
Daejeon, South Korea
dbduscjf@etri.re.kr

Wooyoung Choi
Network Research Division
Electronics and Telecommunications Research Institute (ETRI)
Daejeon, South Korea
wychoi@etri.re.kr

Taekyu Kang
Network Research Division
Electronics and Telecommunications Research Institute (ETRI)
Daejeon, South Korea
tkkang@etri.re.kr

Hogeon Kim
Network Research Division
Electronics and Telecommunications Research Institute (ETRI)
Daejeon, South Korea
steveho.kim@etri.re.kr

Taesik Cheung
Network Research Division
Electronics and Telecommunications Research Institute (ETRI)
Daejeon, South Korea
cts@etri.re.kr

***Abstract*—6G networks aim to achieve ultra-low latency, high reliability, and deterministic communication capabilities crucial for mission-critical applications such as industrial automation, autonomous vehicles, and the tactile Internet. Deterministic Networking (DetNet) provides a promising framework for ensuring bounded end-to-end latency and low jitter. The IETF draft suggests an on-time forwarding mechanism based on Push-In First-Out (PIFO) queues, which eliminates the need for strict synchronization across the entire network. This paper expands on this idea by introducing Resource-based On-time Forwarding (ROF), where a domain controller allocates specific delay budgets for each node to guarantee deterministic service levels. We experimentally validated ROF over long-haul fiber networks spanning up to 1,190 kilometers, achieving sub-15 microsecond jitter and precise latency equalization across multiple network paths. Our results indicate that PIFO-based ROF architectures can meet the deterministic requirements anticipated for future 6G networks.**

Keywords—Deterministic Networking, On-Time Forwarding, PIFO, Bounded Latency, Latency Equalization, 6G

I. Introduction

6G is projected to revolutionize communication systems by offering unprecedented levels of reliability, low latency, and deterministic service guarantees. Applications such as collaborative robotics, remote surgery, and Extended Reality (XR) impose stringent demands on network latency and jitter — often in the order of sub-millisecond bounds and microsecond-level variations[1],[2]. Achieving these requirements over long-haul transport networks remains an open challenge.

IEEE 802.1 Time-Sensitive Networking (TSN)[5] has provided solutions for bounded latency within localized environments but relies heavily on global time synchronization using protocols like IEEE 1588 PTP. Such approaches become impractical at 6G scales, especially across high-capacity optical backbone networks.

IETF Deterministic Networking (DetNet) Working Group [6] introduces on-time forwarding, where packets carry metadata that defines precise transmission windows. Nodes enforce these constraints using programmable scheduling, notably Push-In First-Out (PIFO) queues, thus avoiding synchronization dependencies.

This paper introduces Resource-based On-time Forwarding (ROF) as an architectural enhancement, integrating a centralized domain controller to compute per-node delay budgets. We present the first empirical validation of ROF in a realistic ROF transport scenario, demonstrating its suitability as a foundational technology for future 6G deterministic networking.

II. Background

A. Deterministic Networking in 6G

The sixth-generation (6G) mobile network envisions the seamless convergence of Hyper Reliable and Low Latency Communication (HRLLC), immersive communication, and massive communication to support time-critical applications, such as teleoperation, industrial automation, cooperative autonomous driving, and holographic communication. Among these, support for deterministic services — where packet delivery must meet strict upper and lower latency bounds with minimal jitter and loss — is a key differentiator in 6G over prior generations[3].

To meet these goals, DetNet plays a pivotal role in providing guaranteed End-to-End (E2E) performance, even across multi-domain, multi-hop, and heterogeneous transport environments such as long-haul optical networks. DetNet is a set of mechanisms defined by the IETF to enable bounded latency, extremely low packet loss, and in-order packet delivery for unicast and multicast traffic. It builds upon the principles of TSN, extending them to routed IP/MPLS domains with the following core functionalities[7],[8]:

- Bounded latency provisioning using pre-allocated resources such as bandwidth and buffers along with path-level latency budget enforcement.
- Flow identification and isolation, enabling precise delay tracking and scheduling at each hop.
- Time-aware scheduling through deterministic queuing models such as the PIFO queue, which enables packets to be sorted by a nominal departure timestamp and dispatched within a bounded window defined by upper and lower latency thresholds.
- Packet replication and elimination (PREOF) for reliability, where duplicate packets are forwarded

along disjoint paths and recombined at the destination to eliminate losses.

In 6G architectures, DetNet will integrate with AI-native control planes**,** flexible functional splits**,** and software-defined optical transport**,** making it viable to implement per-flow latency guarantees across dynamically reconfigurable infrastructures. Moreover, synchronization precision at the sub-microsecond level, although not always strictly required, enhances DetNet's ability to measure residence times and enforce time-aligned scheduling[4].

Use cases for DetNet in 6G include:

- Industrial Closed-loop Control: Requires sub-100 μs E2E latency with jitter bounded below 10 μs for precise actuation and sensing cycles.
- Tactile Internet and XR services: Depend on consistent latency and minimal variation to preserve user-perceived fidelity.
- Cooperative Mobility (V2X): Demands deterministic data exchange across multiple moving entities with redundancy and synchronization.

To support these services over IP/MPLS and long-haul optical infrastructure, PIFO-based queueing becomes a practical scheduling mechanism that aligns packet release with controller-defined nominal departure times[10],[11]. As shown in this study, the use of PIFO in conjunction with node-local latency bounding and end-to-end resource orchestration enables on-time packet delivery without requiring per-node flow state or tight time synchronization, thus improving scalability.

Consequently, the integration of DetNet into 6G transport networks, especially when reinforced with experimental validation across geographically distributed testbeds, lays the groundwork for practical deployment of bounded-latency deterministic services at scale.

B. In-time vs. On-time

Fig.1 illustrates the architectural and operational distinctions between in-time and on-time networking. In-time networking aims to forward packets as rapidly as possible, with the E2E latency determined by the sum of the fixed signal propagation delay and variable queuing delays at each network node. Queuing delays are influenced by network congestion and are the primary source of jitter.

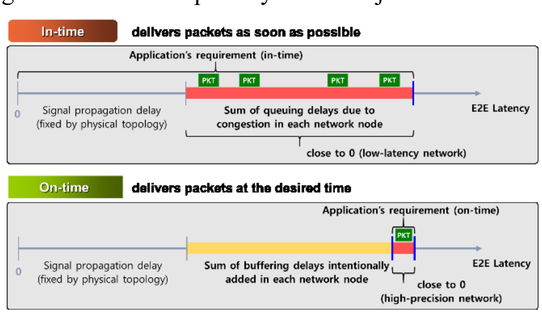

Fig. 1. Comparison of in-time and on-time packet forwarding

Conversely, on-time networking is designed to deliver packets precisely at an application-specified time. Each network node computes the minimum, nominal, and maximum transmission times for every packet based on node-specific lower and upper delay bounds, calculated by the controller. Packets are then scheduled in a PIFO queue, ranked by nominal departure time. While physical signal propagation delays remain unavoidable, intentional buffering delays at nodes allow precise control over packet release times, significantly minimizing jitter.

Moreover, in on-time networks, packets carry fields indicating remaining end-to-end latency bounds. Each node updates these values and uses them to enforce deterministic forwarding behaviors. This approach aligns closely with mechanisms proposed in [9], representing a key technique for achieving bounded latency and low jitter in DetNet environments.

III. ROF Architecture

To ensure deterministic E2E packet delivery in long-haul 6G networks, this study adopts the ROF mechanism as defined in the IETF DetNet framework. ROF is a scheduling paradigm that assigns node-specific buffering budgets along a path and coordinates them using a centralized domain controller, thereby enabling precise packet delivery timing that meets the application's service level objectives (SLOs).

As illustrated in Fig. 2, a network path between a source and destination comprises multiple intermediate nodes N_1, N_2, ..., N_n. Each flow *f* is associated with a latency requirement defined by its lower (SLO_LB(*f*)) and upper bound (SLO_UB(*f*)). The domain controller, which has a global view of network topology and flow requirements, partitions these latency constraints into per-node delay budgets.

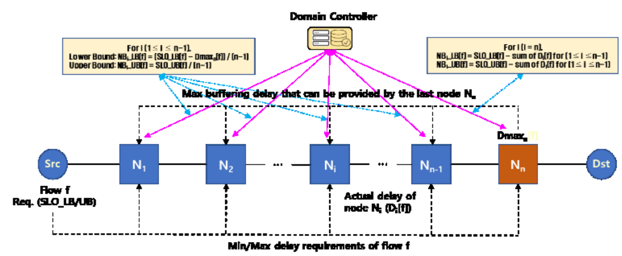

Fig. 2. ROF concept for end-to-end on-time packet forwarding using per-node budgets and centralized domain control

A. Budget Allocation Mechanism

For intermediate nodes N_1 to N_{n-1}, the controller assigns node delay budgets(NB) using the following formula:

$$NB_i^{LB}(f) = \frac{SLO\_LB(f) - Dmax_n(f)}{n-1} \quad (1)$$

$$NB_i^{UB}(f) = \frac{SLO\_UB(f)}{n-1} \quad (2)$$

Where $Dmax_n$ is the maximum latency that can be allocated to the last node N_n.

For the last node N_n, the node delay budgets are adaptively computed to compensate for deviations observed along the previous nodes. Specifically:

$$NB_n^{LB}(f) = SLO\_LB(f) - \sum_{i=1}^{n-1} D_i(f) \quad (3)$$

$$NB_n^{UB}(f) = SLO\_UB(f) - \sum_{i=1}^{n-1} D_i(f) \quad (4)$$

Here, $D_i(f)$ denotes the actual forwarding delay observed at node N_i for flow *f*.

B. Domain Controller Operation

The domain controller performs the following key functions:

- Collects and maintains SLO requirements per flow.
- Monitors node-level delays across the entire path.
- Dynamically distributes buffering delay budgets to meet SLO bounds.
- Ensures on-time delivery even under variable traffic conditions by adjusting the final node's budget as necessary.

This centralized control mechanism aligns with the goals of 6G deterministic networking, delivering bounded latency and jitter while allowing scalability across long-distance optical transport networks.

IV. IMPLEMENTATION

To validate the feasibility of ROF in long-haul networks, a 6G High-Precision Network (HPN) evaluation platform was developed, as illustrated in Fig. 3. This platform is designed to meet the stringent requirements of 6G deterministic services such as HRLLC, which demand ultra-low latency, minimal jitter, and guaranteed end-to-end packet delivery timing.

The 6G-HPN evaluation platform delivers an aggregate system capacity of up to 16Tbps and supports a comprehensive set of line interfaces, including 10GbE, 100GbE, 200GbE, 400GbE, and FlexE/FlexO standards, enabling interoperability across diverse high-speed transport environments. At the heart of this platform is a programmable packet processing implemented on Field Programmable Gate Array (FPGA) based line cards. These line cards integrate deterministic packet forwarding mechanisms compliant with both IEEE TSN (TAS, CQF, ATS, FRER) and IETF DetNet (ROF, PREOF, etc.) standards. The design supports PIFO-based deterministic scheduling, interoperable segment routing, and time synchronization mechanisms based on 1PPS/ToD interface.

Each line card contains two dedicated packet processing modules implemented on an FPGA, which manage packet scheduling, classification, lookup, and deterministic queueing using programmable logic. The packet processing module is an upgraded version of the previously implemented 100G DetNet packet forwarding engine[12], enhanced to support on-time forwarding functions through programmable delay control. These modules are used to evaluate deterministic forwarding schemes under various on-time constraints in long-distance transmission environments. The seamless redundancy block supports TSN/DetNet mechanisms such as Frame Replication and Elimination for Reliability (FRER) and PREOF to enhance end-to-end delivery assurance further.

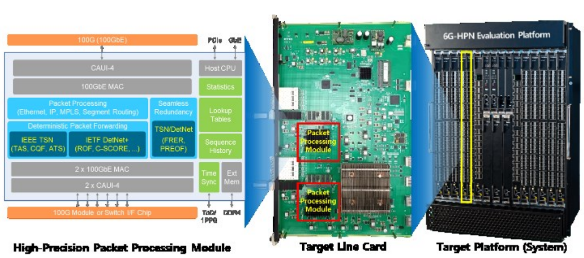

Fig. 3. Architecture and hardware configuration of the 6G high-precision network evaluation platform

V. EXPERIMENTAL SETUP

A. Testbed Configuration

To validate the proposed ROF mechanism under realistic long-haul transport network conditions, a geographically distributed testbed was deployed across three major cities in South Korea—Seoul, Daejeon, and Busan—using the Korea Advanced Research Network (KOREN) optical backbone operating at 100 Gbps.

The end-to-end round-trip path extends across a total of 1,190 km, comprising two optical segments:

- Seoul–Daejeon: Approximately 250 km
- Daejeon–Busan: Approximately 345 km

This configuration emulates a multi-hop optical transport scenario**,** enabling realistic latency and jitter measurements across distances comparable to those of real-world national backbone networks.

As depicted in Fig. 4, the testbed supports several round-trip configurations to evaluate different on-time delivery scenarios:

- SU–DJ RT: Seoul–Daejeon Round Trip
- DJ–BS RT: Daejeon–Busan Round Trip
- SU–DJ x2 RT: Seoul–Daejeon–Seoul–Daejeon Round Trip (looped)
- SU–DJ–BS RT: Full Seoul–Daejeon–Busan Round Trip

Each site is equipped with a deterministic packet processing node, integrated with PIFO queueing and on-time scheduling modules implemented in an FPGA. These nodes are interconnected via high-capacity optical links, supporting 100GbE interfaces as part of the 6G HPN evaluation platform.

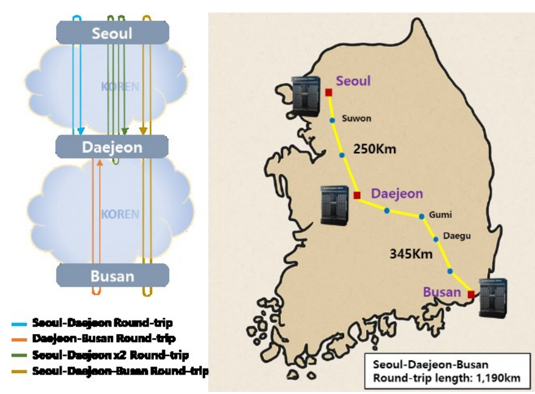

Fig. 4. Experimental testbed topology spanning Seoul, Daejeon, and Busan

B. Test Scenarios

In order to evaluate the performance of on-time forwarding based on deterministic network configuration and PIFO-based scheduling, we designed two separate experimental scenarios.

1) Scenario 1: On-time Target Variation under Single Path

This scenario examines how varying on-time latency requirements affect E2E latency and jitter performance on a single physical path (Seoul-Daejeon-Busan Round Trip). The total round-trip distance is approximately 1,190 km, with an intermediate Reconfigurable Optical Add-Drop Multiplexer

(ROADM) enabling bi-directional optical signal switching and a domain controller managing the delay budgets per flow.

Three types of test packets were used: Best-Effort (BE) traffic for background flow, in-time deterministic networking (DN) traffic, and on-time DN traffic. The on-time traffic flows (flow1–flow5) were assigned different SLO bounds ranging from 6.5ms to 13.5ms. Corresponding delay budgets for each node and the last node were computed and applied by the model, as summarized in Table I.

TABLE I. SLO AND NODE DELAY BUDGET PARAMETERS PER FLOW

Flow No.	SLO_UB	SLO_LB	NB_UB	NB_LB
1	6.51ms	6.5ms	52.9μs	51.4μs
2	7.51ms	7.5ms	195.7μs	193.4μs
3	8.51ms	8.5ms	338.6μs	337.1μs
4	9.51ms	9.5ms	481.4μs	480.0μs
5	12.51ms	13.5ms	1052.9μs	1051.4μs

C. Scenario 2: Path-based Latency and Jitter Analysis

In this scenario, we assess the variation of E2E latency and jitter across four different network paths, each having distinct physical lengths and hop counts, as summarized in Table II. The service requirement was fixed at 6.5ms for all paths, simulating a uniform on-time service requirement. This allows for an evaluation of the scalability and adaptability of the forwarding scheme under a consistent constraint.

TABLE II. SLO AND DISTANCES PER PATH

Path	SLO_UB	SLO_LB	Hop Count	Distance
SU-DJ RT	6.51ms	6.5ms	4	500Km
DJ-BS RT	6.51ms	6.5ms	4	690Km
SU-DJ x2 RT	6.51ms	6.5ms	8	1,000Km
SU-DJ-BS RT	6.51ms	6.5ms	7	1,190Km

VI. EXPERIMENTAL RESULTS

A. E2E Latency and Jitter Comparison under Variable On-Time Requirements

To evaluate the performance of the proposed ROF mechanism, we conducted measurements of E2E latency and jitter under various on-time SLO values, ranging from in-time operation (no added delay) to on-time forwarding with target latencies of 6.5ms to 13.5ms.

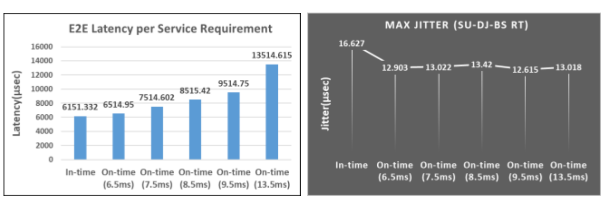

Fig. 5. E2E latency and jitter under variable on-time requirements

As shown in Fig. 5(a), the in-time traffic measured over the full Seoul–Daejeon–Busan round-trip path exhibits an E2E latency of 6,151μs, reflecting the raw propagation and processing delay without any intentional buffering. In contrast, for the on-time flows, the system precisely aligns the packet forwarding time with the configured SLO values. Each on-time flow demonstrates accurate latency alignment, achieving exactly 6.5ms, 7.5ms, 8.5ms, 9.5ms, and 13.5ms, respectively.

Fig. 5(b) shows the maximum jitter (Max–Min latency spread) observed for each service type. The in-time flow exhibited a jitter of 16.627μs, while on-time flows demonstrated reduced jitter values ranging from 12.903μs to 13.420μs, regardless of the increase in SLO values. This result confirms that the ROF mechanism not only achieves precise latency control but also suppresses temporal variability through distributed buffering and deterministic scheduling.

These results confirm that the ROF-enabled forwarding engine can consistently enforce deterministic delays across geographically distributed nodes. The programmable buffering and PIFO-based scheduling modules successfully delay packets to meet the configured timing targets, without exceeding the margin of error, thereby validating the platform's suitability for HRLLC services in future 6G networks.

B. E2E Latency and Jitter Comparison under different paths with a Single On-time Requirement

To assess the robustness of the proposed ROF mechanism across diverse transport topologies, we evaluated end-to-end latency and jitter performance under a fixed on-time requirement of 6.5ms, measured across four different round-trip paths configured on the national optical testbed.

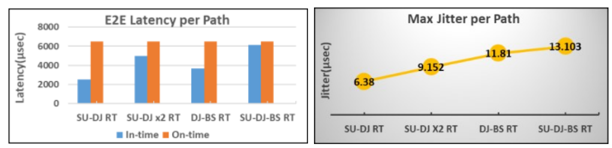

Fig. 6. E2E latency and jitter under different paths

Fig. 6(a) illustrates the comparison of E2E latency measured across four distinct optical paths: SU–DJ RT, SU–DJ x2 RT, DJ–BS RT, and SU–DJ–BS RT. For each path, two service modes were tested: conventional in-time forwarding and deterministic on-time forwarding with a fixed 6.5ms SLO.

In the in-time case, the E2E latency varies significantly depending on the path length and number of hops, reflecting only the inherent propagation and congestion delays. For instance, SU–DJ RT showed the shortest latency (~2.5ms), whereas SU–DJ–BS RT, representing the full round-trip, exhibited a latency of approximately 6.1ms.

Conversely, under on-time forwarding mode, all flows are precisely delayed to align with the 6.5ms target, regardless of the number of hops and geographical distance of the transmission path.

The results verify that the ROF framework correctly allocates per-node buffering budgets, compensating for topological asymmetries and enabling consistent E2E latency across all paths.

In Fig. 6(b), the maximum jitter values of the on-time forwarding mode are compared. While jitter slightly increased with longer paths and additional hops, the values remained well within the tolerable limit for HRLLC services in 6G, which typically allow up to 20μs of jitter. For example, jitter remained consistently below 13μs across all routes, confirming the platform's capability to maintain deterministic timing even under topologically diverse and geographically separated test conditions.

These results underscore the feasibility of scalable deterministic delivery using ROF and PIFO-based scheduling in real-world long-haul optical transport networks.

VII. Conclusion

This paper presents a comprehensive experimental validation of ROF over long-haul optical fiber networks, demonstrating its potential as a deterministic transport solution for 6G networks. Across 1,190 km deployments, we achieved sub-15μs jitter, strict compliance with on-time delivery requirements, and successful latency equalization across paths of varying physical lengths. These results establish ROF, combined with PIFO scheduling, as a viable architecture for enabling 6G deterministic networking services, bridging the gap between theoretical proposals and practical implementations.

Acknowledgment

This work was supported by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT). (2021-0-00715, Development of End-to-End Ultra-high Precision Network Technologies)

References

[1] Wang, C. X. et al, “On the road to 6G: visions, requirements, key technologies, and testbeds”, IEEE Commun. Surv. Tutor. 25, pp. 905–974, 2023.

[2] M. A. Uusitalo et al., “6G vision, value, use cases and technologies from European 6G flagship project Hexa-X,” IEEE Access, vol. 9, pp. 160004–160020, 2021.

[3] ITU-R, ″Framework and overall objectives of the future development of IMT for 2030 and beyond,″ ITU-R M.2441-0, Oct. 2022.

[4] G. P. Sharma, et al., “Toward deterministic communications in 6G networks: State of the art, open challenges and the way forward,” IEEE Access, vol. 11, pp. 106898–106923, 2023.

[5] IEEE Time-Sensitive Networking (TSN) Task Group website, https://1.ieee802.org/tsn/.

[6] IETF Deterministic Networking (DetNet) Working Group website, https://datatracker.ietf.org/wg/detnet/documents/.

[7] N. Finn et al., “Deterministic Networking Architecture,” RFC 8655, Oct. 2019.

[8] N. Finn et al., “Deterministic Networking (DetNet) Bounded Latency,” RFC 9320, Nov. 2022.

[9] Y. Ryoo et al., “On-Time Forwarding with PIFO Queue,” IETF Internet-Draft, draft-ryoo-detnet-ontime-forwarding-02, Feb. 2025.

[10] A. Sivaraman, “Programmable packet scheduling at line rate,” in Proc. ACM SIGCOMM Conf., 2016, pp. 44–57.

[11] A. G. Alcoz et al., “SP-PIFO: Approximating Push-In First-Out Behaviors using Strict-Priority Queues,” in 17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20), 2020.

[12] C. Choi et al., “Implementation and Performance Evaluation of 100Gb/s DetNet Packet Forwarding Engine”, In 13th ICTC, 2022, pp. 1013–1017.