DetNet Working Group
Internet-Draft
Intended status: Informational
Expires: 27 April 2023

J. Joung
Sangmyung University
J. Ryoo
T. Cheung
ETRI
Y. Li
Huawei
P. Liu
China Mobile

24 October 2022

# Asynchronous Deterministic Networking Framework for Large-Scale Networks

draft-joung-detnet-asynch-detnet-framework-01

## Abstract

This document describes an overall framework of Asynchronous Deterministic Networking (ADN) for large-scale networks. It specifies the functional architecture and requirements for providing latency and jitter bounds to high priority traffic, without strict time-synchronization of network nodes.

## Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 27 April 2023.

Table of Contents



## 1. Introduction

Deterministic Networking (DetNet) provides a capability to carry specified unicast or multicast data flows for real-time applications with extremely low data loss rates and bounded latency within a network domain. The architecture of DetNet is defined in RFC 8655 [RFC8655], and the overall framework for DetNet data plane is provided in RFC 8938 [RFC8938]. Various documents on DetNet IP (Internet Protocol) and MPLS (Multi-Protocol Label Switching) data

planes and their interworking with Time-Sensitive Networking (TSN) have been standardized. Technical elements necessary to extend DetNet to a large-scale network spanning multiple administrative domains are identified in [I-D.liu-detnet-large-scale-requirements].

This document considers the problem of guaranteeing both latency upper bounds and jitter upper bounds in large-scale networks with any type of topology, with random dynamic input traffic. The jitter is defined as the latency difference between two packets within a flow, not a difference from a clock signal or from an average latency, as is summarized in RFC 3393 [RFC3393].

In large-scale networks, the end-nodes join and leave, and a large number of flows are dynamically generated and terminated. Achieving satisfactory deterministic performance in such environments would be challenging. The current Internet, which has adopted the DiffServ architecture, has the problem of the burst accumulation and the cyclic dependency, which is mainly due to FIFO queuing and strict priority scheduling. Cyclic dependency is defined as a situation wherein the graph of interference between flow paths has cycles [THOMAS]. The existence of such cyclic dependencies makes the proof of determinism a much more challenging issue and can lead to system instability, that is, unbounded delays [ANDREWS][BOUILLARD]. The Internet architecture does not have an explicit solution for the jitter bound as well. Solving the problem of latency and jitter as a joint optimization problem would be even more difficult.

The basic philosophy behind the framework proposed in this document is to minimize the latency bounds first by taking advantage of the work conserving schedulers with regulators or stateless fair queuing schedulers, and then minimize the jitter bounds by adjusting the packet inter-departure times to reproduce the inter-arrival times, at the boundary of a network. We argue that this is simpler than trying to minimize the latency and the jitter at the same time. The direct benefit of such simplicity is its scalability.

For the first problem of guaranteeing latency bound alone, the IEEE asynchronous traffic shaping (ATS) [IEEE802.1Qcr], the flow-aggregate interleaved regulators (FAIR) [FAIR][Y.3113] frameworks, the port-based flow aggregate regulators (PFAR) [ADN], and the work-conserving stateless core fair queuing (C-SCORE) are proposed as solutions. The key component of the ATS and the FAIR frameworks is the interleaved regulator (IR)), which is described in [I-D.ietf-detnet-bounded-latency]. The IR has a single queue for all flows of the same class from the same input port. The head of the queue (HOQ) is examined if the packet is eligible to exit the regulator. To decide whether it is eligible, the IR must maintain the individual flow states. The key component of the PFAR framework

is the regulators for flow aggregates (FA) per port per class, which regulates the FA based on the sum of average rates and the sum of maximum bursts of the flows that belong to the FA. In the meantime, the key component of the C-SCORE is the packet state that is carried as meta-data. The C-SCORE does not need to maintain flow states at core nodes, yet it is one of the fair queuing schedulers. The service order of the packet is directly inferred from the packet state. It can be implemented based on per-input port FIFO queues. The meta-data to be carried in the packet header is simple and can be updated during the stay in the queue or before joining the queue.

For the second problem of guaranteeing jitter bound, it is necessary to assume that the first problem is solved, that is, the network guarantees latency bounds. Furthermore, the network is required to specify the value of the latency bound for a flow. The end systems at the network boundary, or at the source and destination nodes, then can adjust the inter-departure times of packets, such that they are similar to their inter-arrival times. In order to identify the inter-arrival times at the destination node, or at the network edge near the destination, the packets are required to specify their arrival times, according to the clock at the source, or the network edge near the source. The clocks are not required to be time-synchronized with any other clocks in a network. In order to avoid a possible error due to a clock drift between a source and a destination, they are recommended to be frequency-synchronized.

In this document, strict time-synchronization among network nodes is avoided. It is not easily achievable, especially over a large area network or across multiple DetNet domains. Asynchronous solutions suggested in this document can provide satisfactory latency bounds with careful design without complex pre-computation, configuration, and hardware support usually necessary for time synchronization.

## 2. Terminology

### 2.1. Terms Used in This Document

### 2.2. Abbreviations

## 3. Conventions Used in This Document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

# 4. Framework for Latency Guarantee

## 4.1. Problem Statement

In Section 4, we assume there are only two classes of traffic. The high priority traffic requires latency upper bound guarantee. All the other traffic is considered to be the low priority traffic, and be completely preempted by the high priority traffic. High priority (HP) traffic is our only focus.

It is well understood that the necessary conditions for a flow to have a bounded latency inside a network, are that;

* a flow entering a network conforms to a prescribed traffic specification (TSpec), including the arrival rate and the maximum burst size, and

* all the network nodes serve the flow with a service rate which are greater than or equal to the arrival rate.

These conditions make the resource reservation and the admission control mandatory. These two functions are considered given and out of scope of this document.

Here, the notion of arrival and service rates represent sustainable or average values. A short-term discrepancy between these two rates contributes to the burst size increment, which can be accumulated as the flow passes through the downstream nodes. This results in an increase in the latency bound. Therefore, the value of accumulated burst size is a critical performance metric.

The queuing and scheduling of a flow plays a key role in deciding the accumulated burst size. Ideally, the flows can be queued in separate queues and the queues are scheduled according to the flow rates. In this case a flow can be considered protected. With practical fair schedulers, such as the Deficit Round Robin (DRR), a protected flow still can be affected by the other flows as much as their maximum packet lengths.

If we adopt a separate queue per flow at an output port, and assume identical flows from all the input ports, then the maximum burst size of a flow out of the port, Bout, is given as the following:

$$Bout < Bin + (n-1)L*r/C,$$

where Bout is the outgoing flow's maximum burst size, Bin is the incoming flow's maximum burst size, n is the number of the flows, L is the maximum packet size, r is the average rate of the flow, and C is the link capacity. This approach was taken in the integrated services (IntServ) framework [RFC2212].

The separate queues in the aforementioned case can be too many to be handled in real time, especially at the core of large-scale networks. The common practice therefore is to put all the HP flows in a single queue, and serve them with higher priority than best effort traffic. It is also well known that a proper scheduling scheme, such as the strict priority (SP) scheduling can guarantee service rates larger than the arrival rates, therefore the latency can still be guaranteed. With such a single aggregate queue the flows are not considered protected, however. In this case a flow's burst size in a node can be increased proportionally to the sum of maximum burst sizes of the other flows in the queue. That is,

$$Bout < Bin + (n-1)Bin*r/C.$$

The second product term on the right-hand side represents the amount of increased maximum burst. It is dominated by the term (n-1)Bin, which is the maximum total burst from the other flows.

Moreover, this increased burst affects the other flows' burst size at the next node, and this feedforward can continue indefinitely where a cycle is formed in a network. This phenomenon is called a cyclic dependency of a network. It is argued that the burst accumulation can explode into infinity, therefore the latency is no longer guaranteed.

As such, a flow is required to be protected to a certain level, from the other flows' bursts, such that its burst accumulations are kept within a necessary value. By doing so, the other flows are also protected. The regulators or the fair queuing schedulers are proposed as solutions for such protection in this document. They can decrease the accumulated burst into a desirable level and can protect flows from others. In case of the regulators, however, if the regulation needs a separate queue per flow, then the scalability would be harmed just like the ideal IntServ case. In this document the IR or the regulations on flow aggregates are proposed.

The key requirement for latency guarantee is therefore to have scalability and a certain level of flow protection.

### 4.2. Asynchronous Traffic Shaping (ATS)

The first solution in this document for latency guarantee is the IEEE TSN TG's ATS technology. Essentially it is a combined effort of the flow aggregation per node per input/output ports pair per class, and the interleaved regulator per flow aggregate (FA). The IR examines the HOQ, identifies the flow the packet belongs to, and transfers the packet only when it is eligible according to the initial TSpec of the flow. This solution can have only one queue per FA, but suffers from having to maintain each individual flow state. The detailed description on ATS can be found in [IEEE802.1Qcr].

### 4.3. Flow Aggregate Interleaved Regulators (FAIR)

#### 4.3.1. Overview of the FAIR

In the FAIR framework, the network can be divided into several aggregation domains (ADs). HP flows of the same path within an AD are aggregated into an FA. IRs per FA are implemented at the boundaries of the ADs. An AD can consist of arbitrary number of nodes. The FA can be further subdivided based on the flow requirements and characteristics. For example, only video flows of the same path are aggregated into a single FA.

Figure 1 shows an example architecture of the FAIR framework. The IRs at the AD boundaries suppress the burst accumulations across the ADs with the latency upper bounds intact as they do in IEEE TSN ATS, if the incoming flows are all properly regulated, and the AD guarantees the FIFO property to all the packets in the FA [LEBOUDEC]. It is sufficient to put every FA into a single FIFO queue in a node, in order to maintain the FIFO property within an AD. However, in this case, if cycles are formed, the burst accumulations inside an AD can be accumulated indefinitely. If the topology does not include a cycle and the latency bound requirement is not stringent, then the FIFO queue and the SP scheduler would be allowable. Otherwise, the FAs are recommended to be treated with separated queues and fair-queuing schedulers for flow protection.

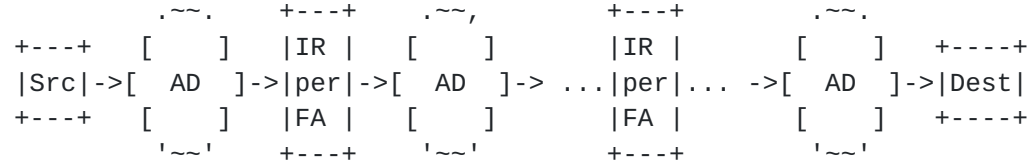


Figure 1: FAIR Framework

### 4.3.2. The performance of the FAIR

The FAIR guarantees an end-to-end delay bound with reduced complexity compared to the traditional flow-based approach. Numerical analysis shows that, with a careful selection of AD size, the FAIR with DRR schedulers yields smaller latency bounds than both the IntServ and the ATS [FAIR].

The ATS can be considered as a special case of the FAIR with the FIFO schedulers, where all the ADs encompass only a single hop. The IntServ can also be considered as an extreme case of the FAIR with fair schedulers and queues per FA, with an AD corresponding to an entire network; therefore, regulators are unnecessary.

## 4.4. Port-based Flow Aggregate Regulators (PFAR)

The IR in the ATS and the FAIR suffers from two major complex tasks; the flow state maintenance and the HOQ lookup to determine the flow to which the packet belongs. Both tasks involve real-time packet processing and queue management. As the number of flows increases, the IR operation may become burdensome as much as the per- flow regulators. Without maintaining individual flow states, however, the flows can be protected to a certain level, as is described in this section.

The ATS and FAIR mitigates the burst increment by placing IRs behind a FIFO system. For example, consider an ATS node with a single queue at an output port for HP traffic. The IR assigned for an input port forms a single queue for the flows from the same input port. Further consider the set of incoming flows from the same input port of the ATS node. Let us call this set of flows the incoming flow aggregate (FAin). If we assume identical FAins from all the input ports, then the maximum burst size of an arbitrary set of flows out of the port, Bout, is given as the following:

$$Bout < Bin + (p-1)B*r/C,$$

where Bin is the sum of maximum burst sizes of the flows within the FAin, B is the sum of initial maximum burst sizes of the flows within the FAin, and p is the number of the ports in the node.

The port-based FA (PFA) is defined as a set of HP flows in the same class, which share the input and output ports in a relay node, such as a switch or router. The only aggregation criteria for a PFA are the ports and the class. The port-based flow aggregate regulators (PFAR) framework puts a regulator for each PFA in an output port module, just before the class-based queuing/scheduling system of the output port module. The PFAR framework sees a PFA as a single flow

with the "PFA-Tspec", {the sum of the maximum initial bursts; and the sum of the initial arrival rates} of the flows that are the elements of the PFA; and regulates the PFA to meet its PFA-Tspec.

The PFARs can be placed at the output port of a node before the output SP scheduler. The architecture is similar to that suggested in the IEEE ATS, except that in the ATS, the IRs are placed instead of the PFARs.

The burst increment of an FA in the PFAR architecture is identical to that in the ATS, which is given as;

$$Bout < Bin + (p-1)B*r/C,$$

where B is again the initial maximum burst size of FAs. However, the regulators in PFAR does introduce additional latency, which is given as

$$D < (Bin - B)/r,$$

where D is the latency within the regulator.

Note that Bout is a function of (n-1)B, not (n-1)Bin; in other words, the burst size out of a node is affected only by the initial burst sizes of the other FAs from different input ports of the node. This property makes the D or Bout do not increase exponentially even in the existence of cyclic dependencies.

With the PFAR, the HOQ flow identification process is unnecessary, and only the PFAs' states, instead of individual flows' states, must be maintained at a node. In this respect, the complexity of process of PFAR is reduced compared to IR of the ATS or the FAIR.

In a recent study [ADN], it was also shown, through a numerical analysis with symmetrical networks with cycles, that PFAR, when implemented at every node, can achieve comparable latency bounds to the IEEE ATS technique.

The ATS, the FAIR, and the PFAR frameworks maintain regulators per FA. The FAs in these frameworks are composed of the flows sharing the same ingress/egress ports of an AD. The ADs can encompass a single hop or multiple hops. The regulators can be the IR or the aggregate regulator. There can be other combinations of AD and regulator type, which could be further investigated and compared to the frameworks introduced in this document.

### 4.5. Work-conserving stateless core fair queuing (C-SCORE)

The generalized processor sharing (GPS) [PAREKH], the weighted fair queuing (WFQ), the virtual clock (VC), and similar other schedulers utilize the concept of finish time (FT) that is the service order assigned to a packet. The packet with the minimum FT in a buffer is served first. We will call these works collectively as the fair queuing (FQ).

As an example, the VC scheduler [ZHANG] defines the FT to be

$$F(p) = \max\{F(p-1), A(p)\} + L(p)/r, \qquad (1)$$

where (p-1) and p are consecutive packets of the flow under observation, A(p)is the arrival time of p, L(p) is the length of p, and r is the flow service rate. The flow index is omitted.

The key idea of the FQ is to calculate the service finish times of packets in an imaginary ideal fluid service model and use them as the service order in the real packet-based scheduler.

While having the excellent flow protection property, the FQ needs to maintain the flow state, F(p-1). For every arriving packet, the flow it belongs to has to be identified and its previous packet's FT should be extracted. As the packet departs, the flow state, F(p), has to be updated as well.

We consider a framework for constructing FTs for packets at core nodes without flow states. In a core node, the following conditions on FTs have to be met.

C1) It has to keep the 'fair distance' of consecutive packets of a flow. That is; Fh(p) >= Fh(p-1) + L(p)/r, where Fh(p) is the F(p) at node h.

C2) The order of FTs and the actual service order, within a flow, have to be kept. That is; Fh(p) > Fh(p-1) and Ch(p) > Ch(p-1), where Ch(p) is the actual service completion time of packet p at node h.

C3) The time lapse at each hop has to be reflected. That is; Fh(p) >= F(h-1)(p), where F(h-1)(p) is the FT of p at the node h-1, the upstream node of h.

C4) The FTs of a flow have to be aligned to the packet arrival times. That is; L(p)/r <= Fh(p)- Ah(p) < Delta.

Delta can be any finite positive value [STILIADIS]. In other words, the Fh(p) should be larger than Ah(p)+L(p)/r, as in (1), yet still should grow at the same rate as Ah(p).

In essence, (1) has to be approximated in core nodes. There can be many possible solutions to meet these conditions. We propose a generic framework for constructing FTs in core nodes, without flow state, in the following.

We denote a 'node' to be an output port of a relay node.

Requirement 1: In the entrance node, it is required to obtain the FTs with (1). That is to obtain F0(p) as in the VC, where 0 denotes the entrance node of the flow under observation.

$$F0(p) = \max\{F0(p-1), A0(p)\}+L(p)/r.$$

Note that F0(p) keeps the fair distances from the FTs of consecutive packets of the flow.

Requirement 2: It is required to increase the FT of a packet by an amount that depends on the node and the packet, dh(p), in a core node h.

$$Fh(p) = F(h-1)(p) + d(h-1)(p).$$

Requirement 3: It is required that dh(p) is a non-decreasing function of p, within a node busy period.

Definition 1: A node busy period is a maximal interval between consecutive node idle periods. During a node idle period, the node has no packet to send.

By Requirements 1, 2, and 3; Conditions 1), 2), and 3) are met.

Requirement 4: It is required that Ah(p)+dh(p) >= A(h+1)(p).

One example of dh(p) is a measured maximum latency of a packet in the node h up until the current packet p, since the start of a node busy period. Let us denote this local maximum latency with uh(p). It may be reset to an initial value during a node idle period. An example of the initial value of uh(p) is the propagation delay from node h to (h+1). By letting dh(p)=uh(p), Requirement 4 is satisfied.

dh(p) may not be a function of p, and dependent only on the node. Then it could be denoted as dh.

One example of dh is letting dh = Uh, where Uh is the latency upper bound in node h for any p. Uh can be a theoretical one, or be obtained by long-term measurements. By letting dh(p)=Uh, Requirement 4 is satisfied.

If Requirement 4 is satisfied then it can be guaranteed Fh(p) >= Ah(p)+L(p)/r, for all h>=0, and it can be proven that Condition 4) is met.

In a core node, the service order of packets from the same input port can be preserved. That is, if Ah(p)> Ah(p') then Ch(p)>Ch(p') for packets p and p' that travel together the nodes (h-1) and h. By preserving the service order of packets from the same input port, using per-input port FIFO queues is possible. An example implementation would be as the follows: The output port module is composed of per-input port FIFO queues. As a packet enters the FIFO queue according to its input port, it should join the queue at the tail and be marked with its FT. The scheduler will examine the smallest FT among the packets at the HoQ of the FIFO queues.

Note that Ah(p)> Ah(p') does not guarantee Fh(p)>Fh(p') when p and p' belong to different flows. For example, p' may have a smaller FT but arrive later while p is in service. However, it is proven that this service completion time discrepancy, C0(p)-F0(p), between real packet system and ideal fluid system is bounded by Lmax/C [PAREKH], where Lmax is the maximum packet length over all the flows, and C is the link capacity.

The meta-data to carry in a packet are Fh(p) and dh(p). These are dynamic and thus need to be updated at every hop. Note that if dh(p) = dh then it can also be signaled out-of-band between the neighboring nodes. Fh(p) can be obtained by a simple summation of two meta-data, and updated during the time interval between the packet arrival and its reaching HoQ of the FIFO queue.

The proposed FT construction framework has advantages of simple FIFO-based implementation and simple meta-data management. We call this solution the work conserving stateless core fair queuing (C-SCORE), which can be compared to the existing non-work conserving scheme [STOICA].

## 5. Framework for Jitter Guarantee

### 5.1. Problem statement

The problem of guaranteeing jitter bounds in arbitrarily sized networks with any type of topology with random dynamic input traffic is considered.

There are several possible solutions to guarantee jitter bounds in packet networks, such as IEEE TSN's cyclic queuing and forwarding (CQF) [IEEE802.1Qch], its asynchronous variations [I-D.yizhou-detnet-ipv6-options-for-cqf-variant], and the latency-based forwarding (LBF) [LBF].

The CQF requires time-synchronization across every node in the network including the source. It is not scalable to a large network with significant propagation delays between the nodes. The asynchronous CQFs are scalable, but they may not satisfy applications' jitter requirements. This is because their jitter bounds cannot be controlled as desired, but are only determined by the cycle time, which should be large enough to accommodate all the traffic to be forwarded.

The systems with slotted operations such as the CQF and its variations turn the problem of packet scheduling into the problem of scheduling flows to fit into slots. The difficulty of such a slot scheduling is a significant drawback in large scale dynamic networks with irregular traffic generations and various propagation delays.

The LBF is a framework of the forwarding action decision based on the flow and packet status, such as the delay budget left for a packet in a node. The LBF does not specify the actions to take according to the status. It suggests a packet slow down or speedup by changing the service order, by pushing packets into any desirable position of a first out queue, as a possible action to take. In essence, by having latency budget information of every packet, the LBF is expected to maintain the latency and jitter within desired bounds. The processing latency required in LBF includes times 1) to lookup the latency budget information on every packet header, 2) to decide the queue position of the packet, 3) to modify the queue linked list, and 4) to update the budget information on the packet upon transmission. This processing latency, however, can affect the scalability especially in high speed core networks.

The ATS, the FAIR, and the PFAR utilize the regulation function to proactively prevent the possible burst accumulation in the downstream nodes. It is not clear whether the LBF can take such preventive action. If so the LBF can also act as a regulator and yield a similar latency bound.

### 5.2. Buffered network (BN)

The BN framework in this document for jitter bound guarantee is composed of

* a network that guarantees latency upper bounds;

* a timestamper for packets with a clock that is not necessarily synchronized with the other nodes, which resides in between, including the source and the network ingress interface; and

* a buffer that can hold the packets for a predetermined interval, which resides in between, including the destination and the network egress interface.

Figure 2 depicts the overall architecture of the BN framework for jitter-bound guarantees [BN]. Only a single flow is depicted between the source and destination in Figure 2. The arrival (an), departure (bn), and buffer-out (cn) times of the nth packet of a flow are denoted. The end-to-end (E2E) latency and the E2E buffered latency are defined as (bn-an) and (cn-an), respectively.

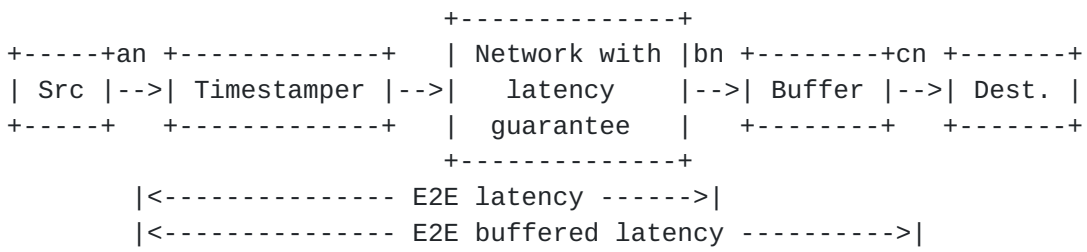


Figure 2: Buffered Network (BN) Framework for Jitter Guarantee

The buffer supports as many as the number of the flows destined for the destination. The destination shown in Figure 2 can be an end station or another deterministic network. The buffer holds packets in a flow according to predefined intervals. The decision of the buffering intervals involves the time-stamp value within each packet.

The network in between the time-stamper and the buffer can be of arbitrarily sized network. The input traffic can be dynamic. It is required that the network be able to guarantee and identify the E2E latency upper bounds of the flows. The network is also required to let the buffer be aware of the E2E latency upper bounds of the flows it has to process. It is recommended that the E2E latency lower bound information is provided by the network as well. The lower bound may be contributed from the transmission and propagation delays within the network.

The time-stamper marks on the packets their arrival times. The time-stamping function can use the real-time transport protocol (RTP) over the user datagram protocol (UDP) or the transmission control protocol (TCP). Either the source or network ingress interface can stamp the packet. In the case where the source stamps, the timestamp value is the packet departure time from the source, which is only a propagation time away from the packet arrival time to the network.

The source and destination do not need to share a synchronized clock. All we need to know is the differences between the time stamps, that is, the information about the inter-arrival times.

### 5.3. Properties of the BN

Let the arrival time of the nth packet of a flow be an. Similarly, let bn be the departure time from the network of the nth packet. Then, a1 and b1 are the arrival and departure times of the first packet of the flow, respectively. The first packet of a flow is defined as the first packet generated by the source, among all the packets that belong to the flow. Further, let cn be the buffer-out time of the nth packet of the flow. Let us define m as the jitter control parameter, which will be described later in detail.

Since buffers can be without cut-through capability, the processing delay within a buffer has to be taken in account. Let gn be the processing delay within the buffer of the nth packet of the flow. The gn includes the time to look up the timestamp and to store/ forward the packet. However, it does not include an intentional buffer-holding interval. By definition, cn - bn >= gn. Let max_n(gn)=g, the maximum processing delay for the flow in the buffer. It is assumed that a buffer can identify the value of g. Let U and W be the latency upper and lower bounds guaranteed to the flow by the network. Let m be the jitter control parameter, W+g <= m.

The rules for the buffer-holding interval decision are given as follows:

* c1=(b1+m-W),

* cn=max{(bn+g), (c1+an-a1)}, for n > 1.

The second rule governing the cn states that a packet should be held in the buffer to make its inter-buffer-out time, (cn-c1), equal to the inter-arrival time, (an-a1). However, when its departure from the network is too late, the inter-buffer-out time should be larger than the inter-arrival time, then hold the packet as much as the maximum processing delay in the buffer, that is, cn=bn+g. The buffer does not need to know the exact values of an or a1. It is sufficient to determine the difference between these values, which can be easily obtained by subtracting the timestamp values of the two packets.

The following theorems holds [ADN].

Theorem 1 (Upper bound of E2E buffered latency). The latency from the packet arrival to the buffer-out times (cn-an), is upper bounded by (U-W+m).

Theorem 2 (Lower bound of E2E buffered latency). The latency from the packet arrival to the buffer-out times (cn-an), is lower bounded by m.

Theorem 3 (Upper bound of jitter). The jitter is upper bounded by max{0, (U+g-m)}.

By setting m=(U+g), we can achieve zero jitter. In this case, the E2E buffered latency bound becomes (2U+g-W), which is roughly twice the E2E latency bound. In contrast, if we set m to its minimum possible value W+g, then the jitter bound becomes (U-W), which is roughly equal to U, while the E2E buffered latency bound becomes U, which is the same as the E2E latency bound.

The parameter m directly controls the holding interval of the first packet. It plays a critical role in determining the jitter and the buffered latency upper bounds of a flow in the BN framework. The larger the m, the smaller the jitter bound, and the larger the latency bound. With a sufficiently large m, we can guarantee zero jitter, at the cost of an increased latency bound.

## 5.4. Frequency synchronization between the source and the buffer

Clock drift refers to phenomena wherein a clock does not run at exactly the same rate as a reference clock. If we do not frequency-synchronize the clocks of different nodes in a network, clock drift is unavoidable. Consequently, jitter occurs owing to the clock frequency difference or clock drift between the source (timestamper) and the buffer. Therefore, it is recommended to frequency-synchronize the source (timestamper) and the buffer.

## 5.5. Omission of the timestamper

For isochronous traffic whose inter-arrival times are well-known fixed values, and the network can preserve the FIFO property for such traffic, then the timestampers can be omitted.

Otherwise the FIFO property cannot be guaranteed, then a sequence number field in the packet header would be enough to replace the timestamper.

## 5.6. Mitigation of the increased E2E buffered latency

The increased E2E buffered latency bound by the proposed framework, from U to almost 2U, can be mitigated by one of the added functionalities given as follows.

1) First, one can measure the E2E latency of a flow's first packet exactly, and buffer it to make its E2E buffered latency be U. Then, by following the rules given in Section 5.3, every subsequent packet will experience the same E2E buffered latency, which is U, with zero jitter. An example of the exact latency measurement may be performed by time-synchronization between the source (timestamper) and the buffer. However, how to measure the latency is for further investigation.

2) Second, one can expedite the first packet's service with a special treatment, to make its latency lower, compared to the other packets of the flow. If we can make the first packet's latency to be a small value d, then every packet will experience the same buffered latency d+U, with zero jitter. Considering that the E2E latency bound is calculated from the worst case in which rare events occur simultaneously, however, the first packet's latency is likely to be far less than what the bound suggests. Therefore, the special treatment to the first packet may be ineffective in real implementations.

### 5.7. Multi-sources single-destination flows' jitter control

The BN framework can also be used for jitter control among multiple sources' flows having a single destination. When a session is composed of more than one sources, physically or virtually separated, the buffer at the boundary can mitigate the latency variations of packets from different sources due to different routes or network treatments. Such a scenario may arise in cases such as

1) that a central unit controls multiple devices for a coordinated execution in smart factories, or

2) multi-user conferencing applications, in which multiple devices/users physically separated can have a difficulty in real-time interactions.

The sources, or the ingress boundary nodes of the network, need to be synchronized with each other in order for the time-stamps from separated sources to be able to identify the absolute arrival times.

## 6. IANA Considerations

There are no IANA actions required by this document.

## 7. Security Considerations

This section will be described later.

## 8. Acknowledgements

## 9. Contributor

## 10. References

### 10.1. Normative References

[I-D.ietf-detnet-bounded-latency]
Finn, N., Boudec, J. L., Mohammadpour, E., Zhang, J., and B. Varga, "DetNet Bounded Latency", Work in Progress, Internet-Draft, draft-ietf-detnet-bounded-latency-10, 8 April 2022, <https://www.ietf.org/archive/id/draft-ietf-detnet-bounded-latency-10.txt>.

[I-D.liu-detnet-large-scale-requirements]
Liu, P., Li, Y., Eckert, T., Xiong, Q., Ryoo, J., Zhu, S., and X. Geng, "Requirements for Large-Scale Deterministic Networks", Work in Progress, Internet-Draft, draft-liu-detnet-large-scale-requirements-05, 20 October 2022, <https://datatracker.ietf.org/api/v1/doc/document/draft-liu-detnet-large-scale-requirements/>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8655] Finn, N., Thubert, P., Varga, B., and J. Farkas, "Deterministic Networking Architecture", RFC 8655, DOI 10.17487/RFC8655, October 2019, <https://www.rfc-editor.org/info/rfc8655>.

[RFC8938] Varga, B., Ed., Farkas, J., Berger, L., Malis, A., and S. Bryant, "Deterministic Networking (DetNet) Data Plane Framework", RFC 8938, DOI 10.17487/RFC8938, November 2020, <https://www.rfc-editor.org/info/rfc8938>.

### 10.2. Informative References

[ADN] Joung, J., Kwon, J., Ryoo, J., and T. Cheung, "Asynchronous Deterministic Network Based on the DiffServ Architecture", IEEE Access, vol. 10, pp. 15068-15083, doi:10.1109/ACCESS.2022.3146398, 2022.

[ANDREWS] Andrews, M., "Instability of FIFO in the permanent sessions model at arbitrarily small network loads", ACM Trans. Algorithms, vol. 5, no. 3, pp. 1-29, doi: 10.1145/1541885.1541894, July 2009.

[BN] Joung, J. and J. Kwon, "Zero jitter for deterministic networks without time-synchronization", IEEE Access, vol. 9, pp. 49398-49414, doi:10.1109/ACCESS.2021.3068515, 2021.

[BOUILLARD] Bouillard, A., Boyer, M., and E. Le Corronc, "Deterministic network calculus: From theory to practical implementation", in Networks and Telecommunications. Hoboken, NJ, USA: Wiley, doi: 10.1002/9781119440284, 2018.

[FAIR] Joung, J., "Framework for delay guarantee in multi-domain networks based on interleaved regulators", Electronics, vol. 9, no. 3, p. 436, doi:10.3390/electronics9030436, March 2020.

[I-D.yizhou-detnet-ipv6-options-for-cqf-variant] Li, Y., Ren, S., Li, G., Yang, F., Ryoo, J., and P. Liu, "IPv6 Options for Cyclic Queuing and Forwarding Variants", Work in Progress, Internet-Draft, draft-yizhou-detnet-ipv6-options-for-cqf-variant-00, 19 June 2022, <https://www.ietf.org/archive/id/draft-yizhou-detnet-ipv6-options-for-cqf-variant-00.txt>.

[IEEE802.1Qch] IEEE, "IEEE Standard for Local and metropolitan area networks -- Bridges and Bridged Networks - Amendment 29: Cyclic Queuing and Forwarding", IEEE 802.1Qch-2017, DOI 10.1109/IEEESTD.2017.7961303, 28 June 2017, <https://doi.org/10.1109/IEEESTD.2017.7961303>.

[IEEE802.1Qcr] IEEE, "IEEE Standard for Local and metropolitan area networks -- Bridges and Bridged Networks - Amendment 34: Asynchronous Traffic Shaping", IEEE 802.1Qcr-2020, DOI 10.1109/IEEESTD.2020.9253013, 6 November 2020, <https://doi.org/10.1109/IEEESTD.2020.9253013>.

[LBF] Clenm, A. and T. Eckert, "High-precision latency forwarding over packet-programmable networks", NOMS 2020 - IEEE/IFIP Network Operations and Management Symposium, April 2020.

[LEBOUDEC] Le Boudec, J., "A theory of traffic regulators for deterministic networks with application to interleaved regulators", IEEE/ACM Trans. Networking, vol. 26, no. 6, pp. 2721-2733, doi:10.1109/TNET.2018.2875191, December 2019.

[PAREKH] Parekh, A. and R. Gallager, "A generalized processor sharing approach to flow control in integrated services networks: the single-node case", IEEE/ACM Trans. Networking, vol. 1, no. 3, pp. 344-357, June 1993.

[RFC2212] Shenker, S., Partridge, C., and R. Guerin, "Specification of Guaranteed Quality of Service", RFC 2212, DOI 10.17487/RFC2212, September 1997, <https://www.rfc-editor.org/info/rfc2212>.

[RFC3393] Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)", RFC 3393, DOI 10.17487/RFC3393, November 2002, <https://www.rfc-editor.org/info/rfc3393>.

[STILIADIS] Stiliadis, D. and A. Anujan, "Rate-proportional servers: A design methodology for fair queueing algorithms", IEEE/ACM Trans. Networking, vol. 6, no. 2, pp. 164-174, 1998.

[STOICA] Stoica, I. and H. Zhang, "Providing guaranteed services without per flow management", ACM SIGCOMM Computer Communication Review, vol. 29, no. 4, pp. 81-94, 1999.

[THOMAS] Thomas, L., Le Boudec, J., and A. Mifdaoui, "On cyclic dependencies and regulators in time-sensitive networks", in Proc. IEEE Real-Time Syst. Symp. (RTSS), York, U.K., pp. 299-311, December 2019.

[Y.3113] International Telecommunication Union, "Framework for Latency Guarantee in Large Scale Networks Including IMT-2020 Network", ITU-T Recommendation Y.3113, February 2021.

[ZHANG] Zhang, L., "Virtual clock: A new traffic control algorithm for packet switching networks", in Proc. ACM symposium on Communications architectures & protocols, pp. 19-29, 1990.

Authors' Addresses

Jinoo Joung
Sangmyung University
Email: jjoung@smu.ac.kr

Jeong-dong Ryoo
ETRI
Email: ryoo@etri.re.kr

Taesik Cheung
ETRI
Email: cts@etri.re.kr

Yizhou Li
Huawei
Email: liyizhou@huawei.com

Peng Liu
China Mobile
Email: liupengyjy@chinamobile.com