In-Network Freshness Control: Trading Throughput for Freshness

Shih-Hao Tseng, SooJean Han, and Adam Wierman

arXiv:2106.02156v1 [cs.NI] 3 Jun 2021

***Abstract*—In addition to traditional concerns such as throughput and latency, freshness is becoming increasingly important. To stay fresh, applications stream status updates among their components, which can congest the network if the update frequency is too high. Tuning to the right frequency is not trivial, especially in the presence of other flows, when network sharing becomes much more involved. Also, sophisticated tuning logic inevitably complicates the design of the endhost devices.**

In this paper, we take an alternative approach. Instead of tuning the update frequency at the end-host, we let the endhost send out updates at its own pace and control the freshness within the network. This In-network Freshness Control (IFC) scheme allows the network operator to improve freshness while providing a fine-grained trade-off with throughput. IFC leverages in-network compute resources to filter out obsolete information during transmission of status updates, while queueing other drop-averse traffic separately to provide high throughput. We provide an analytic study of IFC and then implement IFC as Linux kernel modules. Our experiments show that IFC outperforms existing queueing disciplines by improving both throughput (by up to 40%) and freshness (by up to 50%). IFC can easily be combined with existing methods, e.g., BBR and DCTCP, and is effective even in partial deployments.

***Index Terms*—Age of Information (AoI), in-network traffic control, mixed traffic scheduling.**

I. Introduction

As we step into the era of the Internet of Things (IoT), applications increasingly depend on the ability of the network to deliver information when it is *fresh* so that they can well synchronize their geo-distributed components. Freshness is critical for emerging streaming and IoT applications. For example, Google opens its cloud gaming platform Stadia [1] to consumers across the globe. Live-streaming services such as Facebook Live [2] and YouTube TV [3] are widely used on a daily basis. Microsoft Azure IoT [4] orchestrates interconnected IoT devices while they "observe, identify and understand the world – without the limitations of human-entered data" [5]. The IoT era has also brought a rapid development of smart city technologies, introducing a diversity of safety-critical applications that require consistently available fresh and up-to-date information. Decentralized connected vehicles rely on fresh position estimates at each car/drone to prevent collisions [6], [7], [8], [9], [10], and fresh phasor data updates are crucial for smart grid stabilization [11], [12], [13], [14]. Tasks over large-scale autonomous robotic networks, such as nanosatellite deployment for constellation missions [15] and collaborative data-gathering through swarm-guidance and flocking [16], [17], [18], [19], need fresh measurements to take prompt actions. Based on up-to-date measurements, disaster detection measures like earthquake early warning systems [20], [21], [22] issue timely alarms to save lives. Given these applications, freshness as a performance metric is outpacing more traditional measures like throughput, and freshness-driven flows are expected to occupy a significant share of network traffic in the near future [23].

To stay fresh, the information source of an application needs to refresh its destinations from time to time by streaming the status updates. Ideally, a high update frequency could keep the destinations perfectly in sync. In practice, though, flooding status updates into the network causes congestion that hurts both the synchronization and the throughput of other flows sharing the network. As such, one solution to achieve freshness at the destinations then involves update rate control at the source, which motivates the proposal of the metric *Age of Information (AoI)* [24], [25] and a series of update rate control policies under various settings (see [26] for a survey).

Despite various rate control proposals, very few of them are deployed or even implemented. There are a number of hurdles that have prevented implementation. For example, most analytic work, including the recent system designs [27], [28], focuses on situations with a single class of flows interested in freshness [25], [29], [30], [31], [32], [33], [34], [26]. But in practice, flows are mixed: some require freshness and others prefer throughput, so designs must be able to flexibly handle these varying needs. Especially, it is not clear how endhost-based freshness control can efficiently obtain enough network capacity from other throughput-aggressive legacy flows. Also, a sophisticated rate control policy requires complicated logic at the end-hosts, such as probing, memory, and computation, which might not be available for, say, simple IoT sensors.

Instead of burdening the endhosts, we alternatively explore freshness control within the network – at the intermediate nodes/routers. Notably, for status update streams, outdated data offers no value to the user and can be dropped by the network.[1] Not only does this allow for quicker delivery of fresh information, but it also spares bandwidth for other flows to improve throughput, mitigate congestion, and potentially reduce latency. It is possible thanks to the increasing amount of in-network compute resources in the past decades [36], [37], [38], each network node is now much more capable beyond simple forwarding. Thus, we can leverage the in-network processing capability to filter out obsolete information and

S.-H. Tseng, S. Han, and A. Wierman are with the Division of Engineering and Applied Science, California Institute of Technology, Pasadena, CA 91125 USA (e-mail: {shtseng, soojean, adamw}@caltech.edu).

[1]This opportunity is also noted in [35], in which the authors drop queued old packets at the source, while we study in-network implementation here.

achieve win-win results beyond the capability of traditional networks.

Although promising, there are two major challenges we need to address when filtering obsolete packets in-network. First, the system should operate under practical traffic where flows are mixed – not all flows require freshness. While achieving fresh delivery for status update streams by dropping outdated information, the system should maintain a reasonable throughput level for other drop-averse flows. Additionally, the network must identify outdated packets without incurring too much processing overhead. Dropping stale packets to save transmission time would be meaningless if identification takes longer than transmission. How to devise a low overhead mechanism to enhance freshness is itself a design and implementation challenge.

• **Contributions:** In this paper, we propose *In-network Freshness Control (IFC)*, a design which provides optimization of and balance between freshness and throughput across the network by distinguishing flows based on their performance objectives at each network node and intentionally dropping stale packets to ensure fresh information delivery to endhosts. By performing freshness control within the network, IFC not only simplifies the status update policy at the endhosts, it also naturally addresses network sharing issues. To our knowledge, we are the first to analyze and deal with the problem of mixed freshness and throughput-centric traffic and IFC is the first practical implementation to optimize both freshness (AoI) and throughput.

IFC addresses the challenges mentioned earlier using a per-port architecture. To handle mixed traffic, the system first categorizes flows at each output port into two classes, legacy drop-averse (LDA) flows and AoI flows, and then sorts them into their respective queues governed by different scheduling policies.

A key aspect of IFC is mixing the segregated traffic back together. This procedure addresses the network sharing problem between LDA and AoI flows. IFC does this by prioritizing the two queues alternately according to an *AoI ratio*, which dictates the relative amount of time each queue gets prioritized and allows the operator to explore a whole spectrum of scheduling algorithms to trade off the performance of the two flows. We enforce the AoI ratio through two different schedulers: Time-Division Multiplexing (TDM) and Size-Driven Multiplexing (SDM), where SDM is a special form of fair queueing which prioritizes flows to balance their total size of sent data.

IFC realizes low overhead packet filtering through the adoption of inter-flow First-In First-Out (FIFO), intra-flow Last-In First-Out with preemption only in Waiting (LIFO-W), or IFIL for short, for the AoI queue and a hashed implementation for fast packet preemption. Instead of identifying the age of each packet, IFIL replaces a queued packet with a latter arrived one from the same flow. IFIL is simple to enforce yet its LIFO-W part approximates the outdated packet removal process [29], [26]. Despite its simple criterion, IFIL requires locating the previously enqueued packet, if any, each time a new packet from the same flow arrives. A naive implementation such as linear search through the queue is slow and does not scale. Instead, we propose a hashed architecture to enable fast packet preemption, which ensures the low overhead of IFIL.

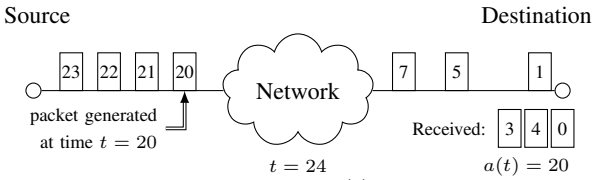

Figure 1: The status update age $a(t)$ at time t is the elapsed time since the freshest update received at the destination at time t was generated at the source. In this example, the current time $t = 24$ and the generated time of each packet is marked accordingly. Upon the arrival of the packet generated at time $t = 1$, the freshest update received at the destination was generated at $t = 4$, and hence the status update age is $a(t) = a(24) = 24 - 4 = 20$ at time $t = 24$. Notice that receiving the outdated (or "out-of-ordered") update (generated at $t = 1$) does not help reduce $a(24)$.

We evaluate the design of IFC using both analysis and implementation experiments. Our analysis gives a bound on the sub-optimality of IFIL. Further, we study properties of schedulers that could potentially outperform both SDM and TDM and argue that they are undesirable in practice.

We implement IFC as Linux kernel modules and evaluate it through emulations. Our results highlight the benefits of designing schedulers to optimize freshness. In particular, IFC incurs small overhead and achieves much shorter AoI and higher throughput than existing queueing disciplines under a wide range network topologies and traffic patterns. With IFC, throughput is boosted up to 40%, and AoI is reduced by up to 50%. Further, IFC can be easily combined with existing queueing disciplines or work with other transport protocols such as BBR [39] and DCTCP [40], improving the performance of these protocols. Finally we show that we can benefit from IFC even when only partial deployment is possible and when the AoI traffic load is light.

In summary, our work presents the first design and implementation that focuses on optimizing freshness (AoI) and throughput of mixed network flows. The design allows the operator to improve freshness while providing a fine-grained trade-off with throughput of traditional flows. This provides a practical demonstration of the ideas around AoI over the past ten years in the information theory community since the works of [24], [25].

II. Background and Motivation

In this section, we first introduce the age of information (AoI) metric for measuring freshness. Then, we describe the opportunities and challenges associated with freshness control. Lastly, we highlight the shortcomings of existing designs when it comes to optimizing freshness.

A. Freshness and the Age of Information

A new metric that has recently emerged to quantify freshness is the *Age of Information (AoI)* [24], which quantifies the degree of synchronization between the information source and

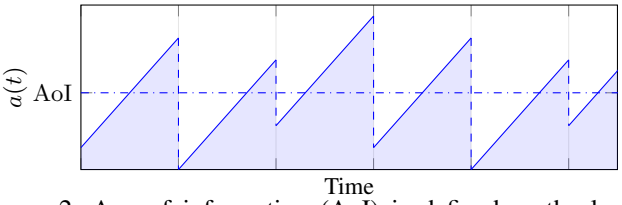

Figure 2: Age of information (AoI) is defined as the long-term time-averaged status update age, which averages $a(t)$ (the shaded area) over time.

its destinations. To synchronize the destinations, the information source streams status update packets through the network. The *status update age* $a(t)$ is then defined as the elapsed time since the freshest update received at the destinations at time t, as illustrated in Figure 1. $a(t)$ reduces upon the receipt of a new update, which results in the saw-like waveform in Figure 2.

Instead of directly referring to the fluctuating $a(t)$, we would quantify the degree of synchronization by the long-term time-averaged performance. Accordingly, we define AoI as the long-term time-averaged status update age, which is given by

$$\text{AoI} = \lim_{T\to\infty} \frac{1}{T} \int_0^T a(t)\ dt. \tag{1}$$

This is illustrated in Figure 2. The existence of the above limit is typically assumed implicitly.

Crucially, AoI/freshness measures something very different from latency. Latency measures how fast a packet can be delivered from the source to the destination, while freshness evaluates how synchronized the source and the destination are. The former is a per-packet property, while the latter depicts a system-wide characteristic. Although shorter latency can potentially help achieve fresher information delivery, freshness does not require a short transmission time for each packet. For example, dropping a packet hurts the latency of transmitting the packet, but it might clear the path for fresher information to arrive at the destination and thus improve the AoI.

B. Opportunities and Challenges

Traditional network designs aim for high throughput and low latency by preventing packet drops. As a result, most existing AoI minimization proposals focus on update rate control to avoid congesting the network, i.e., [24], [10], [34] (see [26] for a survey). Those rate control policies depend heavily on assumptions/measurements of the network/channel properties as well as computation capability at the endhosts, which highly complicate the endhost devices. It is also not clear how those proposals can survive with enough link capacity to update when sharing the network with throughput-hungry LDA flows.

We argue that freshness control need not rely entirely on the endhosts. For the emerging systems/applications such as IoT, the underlying networks are much more capable than simple forwarding. If we allow the network to take a more active role in pursuing fresh delivery, we could reach a "win-win" design between LDA and new AoI flows, such as replacing queued stale packets with fresher ones when received. Controlling freshness within the network also forms a natural platform to share the network: One can separate status update packets from others and pace the former appropriately to maintain freshness without congesting the latter. Techniques such as these have the potential to dramatically improve both freshness and throughput.

To achieve these promising performance gains, we need to address some significant challenges. First, the techniques demand "smarter" involvement of the network. In addition to forwarding, the network needs to sort and process packets based on their properties, which relies on in-network compute resources. Traditionally, it has been hard to incorporate new features like this into the vendor-specific network nodes, and there were few resources available at each node. But with the recent rise of network function virtualization and middleboxes [41], [42], network nodes have become more computationally capable.

Second, the network has to identify stale packets with little overhead. It is not worth recruiting the network to improve freshness if the processing time incurs a much longer delay than transmission. Also, an identification process with large overhead is not scalable. Each network node needs to queue and process a large number of packets, and the identification overhead will become a bottleneck.

Third, not all flows care about freshness; most legacy traffic still prefers high throughput. The coexistence of these two kinds of flows requires an appropriate balance between the pursuit of two different objectives. Also, the behavior of the flows may change depending on how the network treats the different kinds of packets. For instance, if the network always prioritizes status update packets for freshness, it incentivizes the legacy flows to lie about their identity to get prioritized for higher throughput.

C. Insufficiency of Existing Approaches

Overcoming the challenges described in the previous section is the goal of this paper. Before presenting our design, we first highlight why current approaches are not sufficient for optimizing freshness and throughput jointly.

- **Legacy Sensing Systems:** To ensure data freshness, traditional sensing systems need to probe actively [43], [44], [45], reserve specific channels for communication [46], [47], or more aggressively, give up networks entirely and build sophisticated local decision units [9]. These approaches are expensive and energy-hungry. By enabling the network to deliver fresh information, we can greatly simplify end system design and reduce the cost and energy consumption.

- **Priority Queueing:** Priority queueing is a well-known technique for providing different quality of service (QoS) for different flows [48], [49], [50]. However, it only improves the performance of some prioritized flows, while others could potentially suffer unnecessary delay. Further, it does not focus on freshness, although improving delay can sometimes also improve freshness. In contrast, our approach allows the network to actively drop packets, which consequently frees up bandwidth. Thus, we are able to improve the throughput of

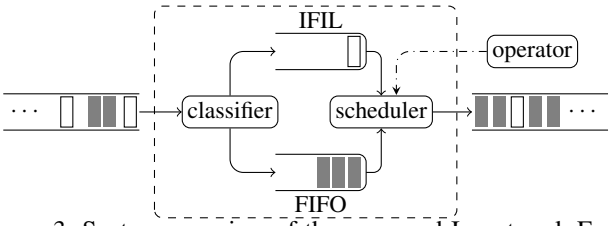

Figure 3: System overview of the proposed In-network Freshness Control (IFC) architecture.

not only the prioritized flows, but also that of all the other flows in addition to the improvements provided for AoI.

• **Flow Completion Time:** The flow completion time problem aims to allocate bandwidth to the flows in order to allow them to send a given amount of data as soon as possible [51], [52], [53]. Although allocating more bandwidth may result in higher throughput, it does not necessarily result in low AoI. Instead, low AoI has more to do with which packet is sent when there is bandwidth available. By not including an intra-flow scheduling component, most solutions to the flow completion time problem miss the opportunity to improve both throughput and AoI.

• **Deadline Scheduling:** Deadline scheduling allocates bandwidth to ensure completion of the flows before the deadlines [54], [55], [56] — a flow that misses its deadline is no longer worth serving. The AoI minimization problem is similar to the deadline scheduling problem in this sense: both problems aim to deliver the "freshest" information, and outdated updates are not worth transmitting. Similar to the flow completion time problem, deadline scheduling improves throughput by dropping flows that are past their deadlines. However, deadline scheduling takes flow-level semantics; packets for a particular flow may be dropped, but doing so would not improve AoI if the entire flow itself is already outdated. Furthermore, for a flow that is still "alive," deadline scheduling does not filter out its old packets at all, something that is crucial for improving AoI.

III. System Overview

In this section, we overview the design of In-network Freshness Control (IFC). IFC does not enforce traffic control at the endhosts. Rather, the endhosts send/update at their own pace, and IFC balances their traffic at in-network forwarding nodes. A diagram of the IFC architecture is shown in Figure 3. Figure 3 shows that incoming traffic, which is a mixture of LDA and AoI packets, is first fed into output ports. Whether a packet is LDA or AoI is declared by its source upon generation.

Notice that a traffic source can mark packets differently to serve its purpose. For example, a status streaming traffic could mark its control packets LDA to ensure a reliable connection, while marking its data packets AoI to achieve fresh delivery. Packet marking can be done by toggling some specific bits in the header (such as the protocol number or flag bits). Although the traffic type is self-claimed, IFC is able to incentivize the participants to stay truthful as shown in Section IV-B.[2]

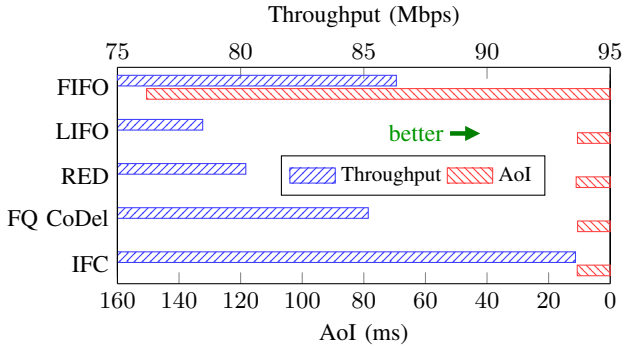

Figure 4: An LDA flow and an AoI flow share one queue, and we measure the throughput (LDA flow) and AoI (AoI flow) under different queueing policies/AQMs. Existing methods (such as FIFO, LIFO, RED, and FQ CoDel) do not distinguish the flows and serving both flows with the same queueing policy/AQM improves one side while hurting the performance of the other. The best solution would be to serve different flows with different queueing policies via IFC, which outperforms all queueing policies/AQMs in achieving both the highest throughput and the lowest AoI.

Since a traffic source could emit both LDA and AoI packets, in this work, we refer to a "flow" (or a "traffic") by the endpoints (processes) and type (LDA/AoI). For instance, a status streaming traffic could maintain two flows, an LDA flow of its control messages and an AoI flow of its data packets. We assume implicitly in this work that each AoI packet is semantically self-contained. At the same time, we remark that the IFC design in this work could be extended to the case where the semantic of one status update spanning over multiple packets.

The simple experiment in Figure 4 suggests that we can improve both throughput and AoI by applying different queueing policies to the different types of traffic. Thus, we keep a *classifier* at each output port that determines the traffic type according to the packet header , then separates the traffic into the two queues with different policies. The standard FIFO queue is used to process LDA packets, while the *inter-flow FIFO, intra-flow LIFO-W (IFIL)* policy governs AoI packets. As its name suggests, different flows are enqueued in FIFO manner, while the packets from the same flow follows LIFO-W policy: Last-arriving packet preempts existing waiting packet in the queue (if any). Operationally, every incoming AoI packet is enqueued at the back of the queue. However, if there exists an AoI packet from the same flow in the queue already, the newly arrived packet preempts (replaces) the existing one. In brief, IFIL schedules different AoI flows in a FIFO manner, which avoids flow starvation, while keeping only the last arrived packet of each flow in the queue.

LIFO-W has been analyzed in [29], [26], where it is referred to as $M/M/1/2*$ and LCFS-W instead. We remark that

[2]An alternative design other than the self-claim scheme would be to identify LDA/AoI packets in-network. But such a design is much more complicated and could prolong the processing time that harms the AoI. As a result, we argue that our self-claim scheme plus truth-revealing IFC design is preferable, given its simplicity and effectiveness.

LIFO-W is different from the preemptive Last-Generated First-Served (LGFS) policy, which outputs the freshest packet of each AoI flow in the queue [31], [57], [33]. To perform LGFS, we need to time-stamp each AoI packet to determine how fresh a packet is, which is costly for both applying and extracting the timestamp information. On the contrary, LIFO-W merely replaces the existing packet by the incoming one, which is much easier to implement. When the packets of the same flow arrive in-order, LIFO-W is a cheap realization of LGFS. Even when the packets arrive out-of-order, LIFO-W can effectively reduce the number of packets from the same flow and clear bandwidth for other flows. As such, we choose LIFO-W over LGFS for a good cost-performance trade-off.

Unlike LIFO-W which keeps only one packet in the queue, IFIL extends LIFO-W by keeping one packet per each distinct flow according to the arrival order. In practice, multiple flows would traverse through and efficient enforcement of IFIL hinges on fast locating the previously enqueued packet. In Section VI, we achieve low overhead IFIL by our hashed implementation.

Note that the trade-off between AoI and throughput depends on the frequency of each type of traffic: if LDA packets arrived at a much higher rate, AoI packets will suffer long queueing delays and be less fresh. This motivates the need for the *scheduler* module. The processed LDA and AoI packets are input into the scheduler, which interleaves the output packets in a way that appropriately balances throughput and freshness. The details of the scheduler and its algorithms are in Section IV.

IV. Scheduler Design

A core component of IFC is the scheduler that interleaves the LDA and AoI packets in a way that balances freshness and throughput. We refer the reader to Figure 3 for how the scheduler fits into the IFC architecture.

In the following subsections, we discuss how to schedule the two queues. Recall that the AoI queue keeps only the latest packet of each flow (IFIL) while the LDA queue performs FIFO. To contrast the schedulers, we focus on the trade-off between throughput and AoI (freshness). To quantify this, we use a trade-off curve, as illustrated in Figure 5. We say a trade-off curve α provides better trade-off than another curve β if α is "closer" to the origin than β, or $\alpha \leq \beta$ under a well-specified partial order (e.g. larger value means "closer" to the origin), as shown in Figure 5.

A. Two algorithms

The most naive scheduler design is to prioritize either the AoI or the LDA queue. However, this naive design misses the whole spectrum of scheduling policies between the two extremes. To capture the intermediate options, our approach is to schedule the AoI and LDA queues according to an *AoI ratio* $\gamma \in [0, 1]$, predetermined by the network operator. We use γ to smoothly transition between prioritizing AoI flows ($\gamma = 1$) and LDA flows ($\gamma = 0$). We consider two ways of

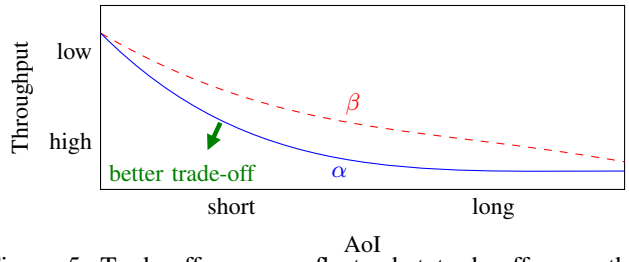

Figure 5: Trade-off curves reflect what trade-offs a method can offer by tuning its parameters. A trade-off curve provides better trade-off if it is closer to the origin. In this case, α provides better trade-off than β.

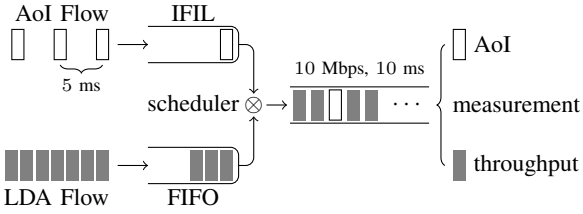

Figure 6: Single link setup for scheduler comparison. The AoI and LDA flows arriving at the node are classified into two queues. The scheduler then schedules the two queues to send through a 10 Mbps link with transmission delay 10 ms. The measurements are then taken at the output of the link.

prioritization: Time-Division Multiplexing (TDM) and Size-Driven Multiplexing (SDM).

• **Time-Division Multiplexing (TDM):** A classical approach for interleaving two queues is time-division multiplexing (TDM). The idea is simple: we partition a given time frame T_f into two intervals, one for each queue. Within each time interval, the corresponding queue is prioritized. Accordingly, we can define γ as

$$\frac{T_{\text{AoI}}}{T_{\text{LDA}}} = \frac{\gamma}{1-\gamma}$$

where T_{AoI} and T_{LDA} denote the total time when the AoI and the LDA queue are prioritized, respectively. The lengths of the AoI and LDA intervals are then set to γT_f and $(1-\gamma)T_f$.

Although TDM is a common and intuitive algorithm, there are two potential issues in our setting. First, since the per-packet scheduling decisions are non-preemptive, a packet might still occupy the link while the next interval starts, which disturbs the ratio between the two intervals. This can be addressed by prolonging the next interval accordingly.

The second issue is that the performance of TDM is tied to the selection of the time frame and the inter-arrival time of the AoI flows. We apply different TDM time frames to schedule the AoI and LDA queues using the setup in Figure 6 and show the results in Figure 7. Notice that the shorter the time frame, the closer the trade-off curve is to the origin, meaning that it yields better performance. However, the curve also becomes less smooth, so there is less guarantee that such a good performance will be consistently maintained.

• **Size-Driven Multiplexing (SDM):** A contrasting approach to TDM is size-driven multiplexing (SDM), which prioritizes the queue based on size. More specifically, let S_{AoI} be the

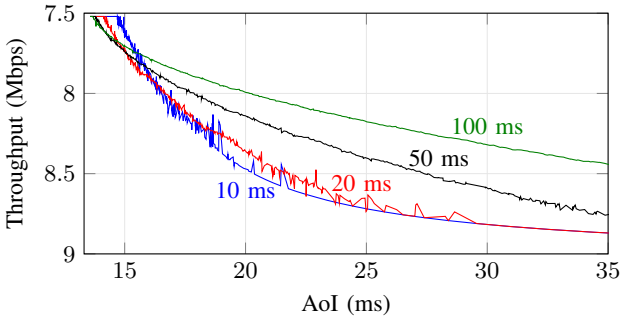

Figure 7: Time-division multiplexing (TDM) trade-off curves under different time frames under the setting in Figure 6. TDM partitions a given time frame into two intervals prioritizing AoI and LDA queues respectively. Different time frame lengths lead to different trade-off curves. A longer time frame gives a smoother trade-off curve, but the throughput is lower and the AoI is longer.

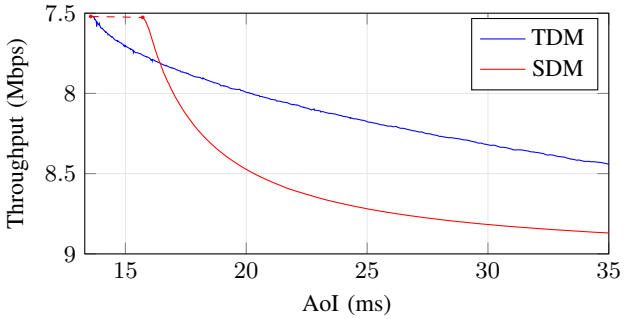

Figure 8: Trade-off curves of different schedulers in the single link system as in Figure 6. TDM (time frame $= 100$ ms) provides better trade-off than SDM when AoI is short. However, SDM outperforms TDM when throughput is high. The trade-off curve of SDM has a discontinuous AoI surge at the upper-left corner, which we call it the transition point effect.

total size of sent AoI packets and S_{LDA} be defined similarly. SDM aims to ensure

$$\frac{S_{\text{AoI}}}{S_{\text{LDA}}} = \frac{\gamma}{1-\gamma}.$$

To do so, we define the budget function as $b(\gamma) = \gamma S_{\text{LDA}} - (1-\gamma)S_{\text{AoI}}$. so that SDM tries to maintain $b(\gamma) = 0$. When $b(\gamma) > 0$, the AoI queue is prioritized, and when $b(\gamma) < 0$ the LDA queue is prioritized. However, prioritizing a queue does not necessarily imply an increase in the service rate of the corresponding kind of packets, especially for the AoI queue. The scheduler can only obtain AoI packets at a rate up to its incoming rate.

Notably, at the steady state, SDM maintains a constant pace of outputting AoI packets, which slows down as γ decreases. This periodic service suffers from an issue known as the *transition point effect*, i.e., the AoI is prolonged significantly when the output period is slightly longer than its inter-arrival time. Therefore, SDM cannot outperform TDM under all situations.

On the other hand, one could tune TDM to realize a given point on the SDM trade-off curve by choosing T_f as the SDM periodic cycle and γT_f the time to transmit one AoI packet. However, we emphasize that TDM cannot replace SDM for the following reasons. First, the trade-off curves depict the performance at the steady state, but LDA and AoI traffic could arrive in bursts, which are handled differently by TDM and SDM. Second, we don't know the SDM cycle a priori, and TDM has to commit to a T_f in advance. Third, TDM requires a timer or the access to system time. In comparison, SDM only needs a counter, which is much simpler and results in much smoother and monotonic trade-off curve. Monotonicity is critical for the operator to tune γ for the most desired performance.

B. Comparison and Discussion

To end this section, we compare and contrast TDM with SDM using the experimental setup shown in Figure 6. In our experiments, the AoI and LDA flows enter a 10 Mbps link of propagation delay 10 ms. The inter-arrival time of the AoI packets is 5 ms and each packet has 1024 bytes. The AoI and throughput are measured at the output of the link for AoI and LDA flows, respectively.

Table I: Truthful Revelation Property. The AoI and LDA flows can act either truthful (T) or deceptive (D). The corresponding AoI (ms) and throughput (Mbps) are measured and compiled in the table. Both being truthful is the Nash equilibrium.

		LDA	
		T	D
AoI	T	(13.57 ms, 7.52 Mbps)	(13.38 ms, 0.41 Mbps)
	D	(131.4 ms, 7.52 Mbps)	(13.38 ms, 0.41 Mbps)

Figure 8 shows the trade-off curves of TDM and SDM. The time frame of the TDM is 100 ms in this figure, which is the smoothest trade-off curve in Figure 7. The figure highlights that TDM leads to better trade-off when AoI is short, while SDM outperforms TDM when throughput is high. The reason is that TDM degrades AoI more gracefully by distributing the fixed time frame for the two queues, while SDM suffers the transition point effect as it prioritizes AoI queue periodically; note that the AoI grows significantly when the period increases slightly above the minimum cycle. On the other hand, as shown in Figure 7, the fixed time frame also impacts the performance of TDM, which makes it underperform at higher throughput.

• **Truthful Revelation:** An important remaining issue is that, to serve AoI and LDA flows differently, we need the flows to identify themselves. This will not happen if a flow benefits from revealing information deceptively. However, both TDM and SDM incentivize the flows to truthfully reveal their identities.

We remark that it is not always true that a scheduling policy differentiating AoI and LDA would incentivize the participating flows to reveal their true identities. Consider the following naive policy: always prioritizing AoI packets and scheduling AoI packets by FIFO while controlling the LDA queue length by RED or FQ CoDel. Without AoI packets, RED or FQ CoDel could improve the throughput of LDA flows. Prioritizing AoI packets over LDA packets can also shorten the AoI. However, in this case, LDA flows have the incentive

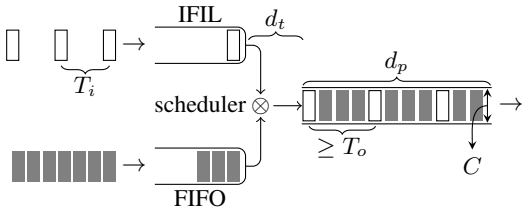

Figure 9: For analysis, we model the network as a link with capacity C and propagation delay d_p. We denote the inter-arrival time of the AoI flow by T_i, the minimum output period by T_o, and the transmission delay by d_t. Modeling LDA flows as fluid, we assume that LDA traffic can fully occupy the spare link bandwidth.

to pretend as AoI flows to get prioritized. Fortunately, TDM and SDM do maintain the truthful revelation property.

To illustrate this truthful revelation property, we consider again the simple setup in Figure 6 and allow the flows to claim their identities freely, i.e., an AoI flow can pretend to be an LDA flow and be scheduled accordingly. The results in Table I show that being truthful is the dominant strategy for the LDA flow. When the LDA flow plays deceptive, its packets get dropped and cleared the path for shorter AoI, which significantly hurts its throughput and hence not preferable. Further, it also follows that the best strategy for the AoI flow is to be truthful.

Besides the empirical example above, we remark that one can show, under some mild assumptions on the queueing policies (i.e., the LDA queue improves throughput and the AoI queue improves AoI), that the truthful revelation property holds for TDM and SDM regardless of input traffic.

V. Analysis

In this section, we take an analytic approach to investigating the trade-offs in the IFC design. We show that IFC, despite its simple design, yields an AoI not too far from the optimum, and derive properties of scheduler designs that are better than the one we propose, but show that there is a price to pay. We defer the proofs of the theorems and the discussions on how to set the AoI ratio to achieve the optimal trade-off between throughput and AoI to Appendices.

• **Model and Notation:** We model the system as a single-link channel with capacity C, as illustrated in Figure 9. An AoI flow of inter-arrival time T_i is queued at the AoI queue, which is governed by IFIL. A scheduler, as described in Section III, allows the AoI queue to send with the minimum output period T_o. Once the AoI queue is scheduled to send, an AoI packet is dequeued. The node serializes the update with transmission delay d_t and sends it through the link with propagation delay d_p. Note that the transmission delay d_t applies only to AoI packets and it is usually proportional to the size of the packet. Here, we simply model it as a constant. Finally, we assume the LDA flow can always fill the spare link bandwidth in between AoI updates.

A. IFIL: Simple yet Effective

IFC uses IFIL to schedule the AoI queue. IFIL is simple to implement and its single flow version, LIFO-W, has been shown effective for Poisson/Bernoulli sources [29], [31]. However, IFIL can also perform quite badly under other arrival processes. For example, consider a pulse train input with inter-arrival time $T_i = 1$ ns and a server that takes time d_t to process each packet. Clearly, if $d_t \leq T_i$, every packet will arrive at an idle link and begin transmission immediately. Suppose $d_t = 1.9$ ns. The second packet will arrive while the first is still being transmitted, and it will incur an additional age of $d_t - T_i = 0.9$ ns while it waits for the link to become available. Under IFIL policy, the resulting AoI is 3.74 ns. Not only is this suboptimal, but a better policy is to simply have the link wait for an additional 0.1 ns before transmitting queued packets. Doing so reduces the AoI to 2.9 ns.[3]

Although this shows IFIL can be suboptimal for pulse trains, it leaves open the question of just how suboptimal it is. In the following theorem, we show that IFIL only has a small constant factor loss compared to the optimum, even in the worst case of pulse trains. Given the simplicity of IFIL, we argue that this provides good justification for the use of it in IFC.

Theorem 1. *Consider an AoI pulse train of inter-arrival time T_i queued at the source going through a link. The AoI governed by IFIL is at most $\frac{T_i}{2}$ more than the AoI governed by the optimal policy.*

B. Price to Pay for Better Schedulers

Figure 8 shows that neither TDM nor SDM uniformly outperforms the other. Thus, a natural question is whether it is possible to design a scheduler that outperforms both TDM and SDM, taking the best from each. Below we derive a property that must be true of any scheduler that is better than both TDM and SDM, then illustrate why one may not want to use such a design.

Theorem 2. *Consider the system in Figure 9. If a scheduling algorithm outperforms TDM under any given d_p and $T_i > 2d_t$, it must have $T_o = 0$ or it is not work conserving for the AoI queue.*

Theorem 2 highlights that if a scheduler outperforms both TDM and SDM, it must impose zero gap between the service of two consecutive AoI packets, or it should be able to wait even when the AoI queue is not empty. We remark that Theorem 2 does not require the scheduler to keep serving the AoI queue. Rather, it says that the scheduler must either occasionally allow intermittent bursts of AoI packets or deliberately put AoI packets on hold. Of course, such bursts come at the direct expense of the LDA packets and so these bursts cannot be too frequent. Likewise, the waiting policy is nontrivial to derive. This highlights the drawbacks of such a design. To enable intermittent bursts of AoI packets or delayed output, the scheduler needs an additional state to keep track of the number of output AoI packets. The introduction of the additional state complicates the scheduler design and slows down the processing of each packet. Further, setting the

[3]This waiting policy is also noted and discussed in [32] for "generate-at-will" sources.

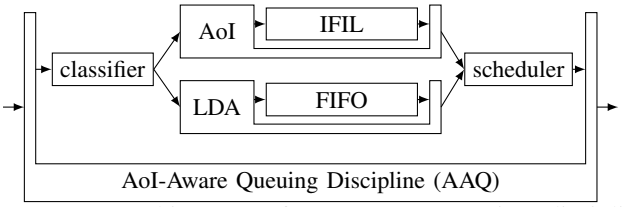

Figure 10: Architecture of AoI-aware queueing discipline (AAQ).

frequency of the AoI bursts/halt period so as to balance the AoI and LDA traffic depends on fine-grained details of the workload, which adds additional complexity and overhead. In contrast, by using either TDM or SDM, as appropriate for the networks, we have a significantly lighter implementation and operation overhead while performing nearly optimally in practice.

VI. Implementation

At the high level, we implement IFC as a set of Linux kernel modules, consisting of two queueing disciplines: AoI-aware queueing (AAQ) and IFIL. We remark that this is not the only way to potentially implement IFC. For instance, one could instead use servers as classifiers and schedulers. We have also successfully implemented IFC in DPDK [58] to demonstrate the potential of a smartNIC deployment. However, DPDK is vendor-specific and its deployment in a virtualized environment (such as cloud-based content-delivery networks) is highly involved. Therefore, we also implement IFC as Linux kernel modules so that it is easy for embedded and commodity Linux systems, such as connected vehicles, networked robotic systems, and public clouds, to import our design without further modification.

• **AoI-Aware Queueing Discipline (AAQ):** As shown in Figure 10, we pack the classifier and the scheduler into one queueing discipline, AAQ. The classifier is a function that categorizes packets based on their header. In our implementation, the classifier distinguishes packets using their layer 4 protocol number as we deem UDP traffic the AoI traffic in the experiments. One can adopt other classifiers according to their definition of AoI traffic. For instance, one can also design a special layer 3 protocol for AoI traffic, and the corresponding classifier filters traffic according to the layer 3 protocol number.

AAQ has two subclasses for AoI and LDA flows. Each subclass handles the traffic through the queueing discipline it obeys. By default, we employ IFIL for AoI flows and FIFO for LDA flows. We can change the queueing disciplines of the subclasses to handle the flows differently. The scheduler can prioritize these two subclasses according to its scheduling policy, e.g., TDM or SDM.

• **IFIL:** To enforce IFIL, outdated packets within the queue need a way of quickly being replaced with fresher arrivals. We could do so by linearly searching through the queue, but this introduces significant overhead. Instead, we store the packets in the queue as a doubly-linked list of their meta-data and maintain a hash table to locate the packet from each flow. The architecture is shown in Figure 11. Since we keep only one packet per flow in the IFIL queue, we can pre-allocate the memory for the hash items to reduce overhead due to memory allocation. We remark that the total number of hash items and meta-data is bounded by the length of the queue. Under moderate queue size, e.g., 10^4 to 10^5 packets, the whole data structure (hash items and meta-data) could easily fit in the cache.

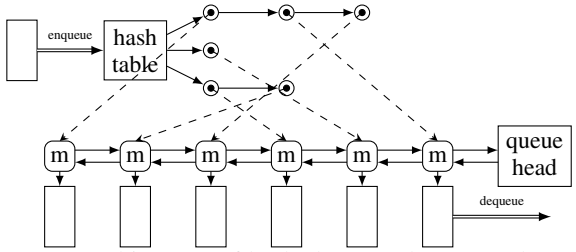

Figure 11: Architecture of hashed IFIL. The rectangles are the packets and the boxes marked by m are the meta-data of the packets.

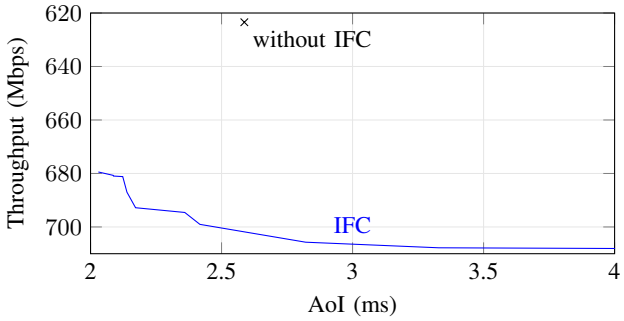

Figure 12: The AoI/throughput trade-off of a one-link (physical, 1 Gbps) example. The results show that our Linux kernel implementation can sustain at least a 1-Gbps link, providing the IFC benefits without incurring significant overhead.

• **Effectiveness:** To demonstrate that our Linux kernel implementation is effective with acceptable overhead, we physically connect two carefully time-synchronized computers by a 1-Gbps link, pump mixed AoI/LDA traffic through it, and measure the AoI/throughput trade-offs in Figure 12. The results show that our Linux kernel implementation can serve at least a 1-Gbps link without incurring significant overhead. We remark that time-synchronization is only needed because we want to measure the AoI accurately. In practice, we don't need the computers to be time-synchronized to perform IFC. To gather accurate AoI measurements over multiple computers, we emulate the network and use 100 Mbps virtual links in Section VII to avoid challenging multi-computer time-synchronization. But given Figure 12, we expect the results extend to physical links of 1 Gbps or higher bandwidth.

VII. Evaluation

We evaluate IFC through extensive emulations, reported on in this section. We begin by describing the setup in Section VII-A and then address the following questions:

- Section VII-B: What is the overhead of adopting IFC? How much does our design reduce the overhead?

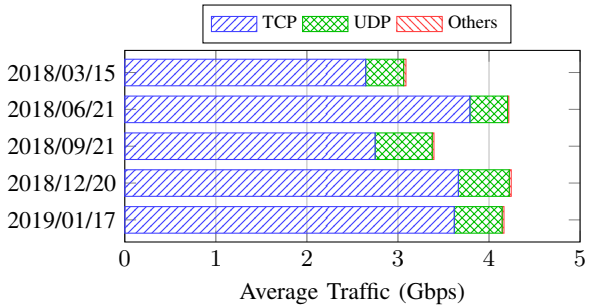

Figure 13: In the CAIDA Internet trace of 2018 and 2019, most of the traffic is TCP, and UDP traffic contributes about 10% of the total.

- Section VII-C: How does IFC perform in comparison to state-of-the-art methods? How does the AoI ratio γ affect the performance of IFC? How well does IFC perform under different traffic patterns?
- Section VII-D: Can IFC be combined with existing methods? How much improvement does this provide?
- Section VII-E: How much does the network benefit from partial IFC adoption?
- Section VII-F: How do the benefits of IFC depend on the load of AoI traffic?

A. Setup

We emulate IFC in Mininet [59]. Mininet allows us to measure AoI using the same system clock. We also consider more practical evaluations over real wide-area networks, but synchronization issues make it hard to gather AoI measurements accurately. We remark that the AoI measurements are essential for performance quantification and discussion, but in practice, freshness difference can be experienced directly and so there is no need to measure it.

• **Network topologies:** We emulate network topologies from Microsoft's SWAN [60], Internet2 [61], and Google's B4 [62], [63].

• **Existing Methods:** In addition to IFC with SDM/TDM, we compare against the existing AQM/TCP proposals including FIFO (First-In First-Out), LIFO (Last-In First-Out), RED [64], FQ-CoDel [65], BBR [39], and DCTCP [40]. We directly use Linux `tc pfifo`, `red`, `fq_codel`, `tcp bbr`, and `dctcp` as their implementations using default Linux parameter settings. For pure AQM methods, we pair them with default Linux TCP (CUBIC).

• **Traffic Characteristics:** We experiment with traces from CAIDA [66]. We report on five traffic traces in 2018 and 2019 from the passive monitor equinix-nyc of [66]. Focusing on one link direction, we plot the results in Figure 13.

Figure 13 shows that UDP traffic contributes about 10% of the total traffic in the latest trace (2019). We also calculate the median of TCP and UDP packet sizes, which are 1452 and 229 bytes respectively. In the emulations, we generate LDA/AoI traffic according to the TCP/UDP ratio in the trace. We first calculate the UDP link capacity portion and assign the interarrival time correspondingly to reach the portion. TCP traffic then performs congestion control to fill the available bandwidth.

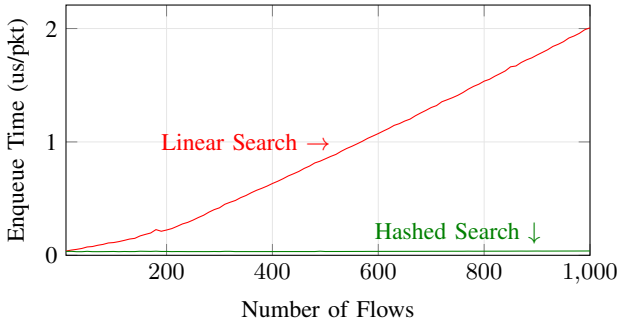

Figure 14: The overhead of IFIL is small with hashed design, about 34.61 ns per packet. In exchange, we get huge AoI improvement.

B. Overhead of IFIL and Schedulers

Since IFC is designed to achieve low AoI, we have to ensure that the overhead is low when adopting it in the system. In particular, we are interested in the overhead of two components: the IFIL queue and the scheduler.

We first examine the overhead of the IFIL queue. When a packet arrives at a IFIL queue, it replaces the packet of the same flow, if there exists one, which requires searching through the queue. The naive method would be to search linearly through the queue, which can be inefficient when the queue stores a large number of flows. Instead, we propose hashed implementation in our design. In Figure 14, we quantify the overhead by pumping packets from a number of flows into the queue and measuring the enqueue time. Figure 14 shows that the proposed hashed implementation significantly outperforms linear search. The overhead of the hashed implementation is around 34.61 ns, while the linear search overhead grows linearly. Notice that IFIL keeps only one packet per flow, meaning that we only need to keep a small number (number of flows) of hashed items and packet meta-data to perform hashing, which could fit it the cache.

Another overhead we need to consider is the scheduler overhead. For comparison, we prepare two fully loaded FIFO queues as the AoI and LDA queues and adopt the scheduler to dequeue packets from the two queues. The scheduler dequeues according to the scheduling policy, and we measure the dequeue overhead in terms of the time. If we keep dequeueing packets from non-empty queues without scheduling, the overhead is 3.79 ns (per packet). On the other hand, the dequeue overhead for SDM and TDM are 7.07 and 21.98 ns, respectively, regardless of the AoI ratio. This validates the fact that both SDM and TDM have very low overhead, though TDM requires more time since it inquiries system time per packet arrival.

C. Comparison Against Existing Methods

• **Trade-off Curves of SDM and TDM:** We start by comparing SDM and TDM with some state-of-the-art queueing disciplines: FIFO, RED, and FQ-CoDel. We emulate four different network topologies in Mininet: single link, SWAN, Internet2, and B4. For single-link, we also include the result of LIFO to verify the intuition that LIFO is beneficial for AoI but significantly harmful for throughput. The capacity of each link is 100 Mbps. We randomly generate the source-destination

pairs (with probability 0.5) and launch one LDA flow and one AoI flow per pair. We run the emulation for 5 minutes and then measure the average throughput of the LDA flows and the average AoI of the AoI flows. The results are shown in Figure 15, where we also draw the trade-off curves formed from the data points of SDM and TDM.

Figure 15 highlights that SDM and TDM can both outperform all state-of-the-art queueing disciplines under most choices of the AoI ratio (γ). It is not surprising: IFC intentionally drops AoI packets which yields the bandwidth for higher throughput and shorter AoI, which directly explains why IFC with FIFO governed LDA queue outperforms FIFO for mixed traffic. In fact, the emulation results suggest that treating LDA and AoI traffic differently (such as IFC) could be more effective than handling mixed traffic with sophisticated type-blind queueing policy. In the results, SDM performs the best when γ is between 0.1 and 0.3, while TDM performs better when $\gamma \in (0.5, 0.9)$. TDM has a wider tuning region, but SDM usually gives better performance.

• **Performance under Random Traffic Patterns:** In Figure 15, IFC outperforms the existing queueing disciplines under most of γ for a specific traffic pattern. In practice, however, we would like to set the γ first, without knowledge of the traffic. So, it is important to examine the performance of IFC under different traffic patterns given a fixed γ.

Figure 15 shows that both SDM and TDM perform well when γ is around 0.5. Therefore, in Figure 16, we run 100 random traffic patterns with $\gamma = 0.5$ for both SDM and TDM. In the figure, each point corresponds to a traffic pattern, and we normalize the measured results by the SDM result for each traffic pattern. The results show that SDM outperforms other algorithms in nearly all cases. There are only few cases in SWAN where FQ-CoDel can outperform SDM in both throughput and AoI. FQ-CoDel performs much worse in Internet2 and B4. SDM achieves 20% more throughput than RED while providing comparable AoI. TDM, on the other hand, sacrifices some AoI in exchange for higher throughput under $\gamma = 0.5$.

D. Combination with Other Methods

A key benefit of IFC is that it can easily be combined with other methods to improve their performance. Specifically, we can replace the FIFO queue in IFC with a different queueing discipline in order to handle LDA flows differently. In Table II, we examine the performance improvements of FIFO, FQ-CoDel, RED, BBR, and DCTCP after combining them with IFC (with scheduler SDM and $\gamma = 0.5$) under different network topologies. For each experiment, we conduct 100 random traffic instances, normalize the measurements by the performance of IFC, and calculate the mean and the standard deviation.

Table II shows that the combination with IFC can improve throughput by 10% to 40% and shrink AoI up to 50%. Notably, standard IFC (which schedules LDA flows using FIFO) achieves comparable performance with some more sophisticated combinations with BBR and DCTCP. Perhaps counterintuitively, the AoI under RED and DCTCP does not improve much. The reason is that they already try to maintain a short queue. However, the throughput does get boosted by 10% to 20% as a result of the combination.

E. Partial Adoption of IFC

In practice, it is typically impossible to upgrade the whole network at once. Instead, new approaches are usually adopted in parts of the network first. For example, for distributed connected vehicles, some of the new vehicles might be capable of performing IFC while the older models cannot. For power systems, the phasor data is routed through several autonomous systems, where some of them might (partially) support IFC. Similarly, enterprise networks could upgrade their routers to adopt IFC in batches instead of at once. In this section, we show that IFC still leads to significant improvement with only partial adoption.

Figure 17 shows the results of emulating the Internet2 network with each link adopting IFC with some probability which, when 0, means that all links obey some legacy queueing discipline (FIFO, RED, or FQ CoDel) and, when 1, means that IFC is fully adopted. The results show that, as we increase the adoption of IFC, the throughput increases and the AoI decreases. When we start with all FIFO or FQ CoDel links, the increasing coverage of IFC links improves both throughput and AoI, while SDM and TDM emphasize on different aspects: SDM leads to shorter AoI and TDM leads to higher throughput. Notably, for the case where the legacy deployment is RED, increasing adoption can improve the throughput without AoI degradation (SDM) or trade a little AoI for much higher throughput (TDM).

F. Benefits under Varying AoI Load

Through the previous experiments, we demonstrate that IFC is helpful for mixed LDA and AoI traffic. Since AoI traffic is not yet included in most traces nowadays, we use the UDP traffic as a proxy for AoI traffic. However, in practice, not all UDP traffic is AoI traffic. Therefore, we conduct experiments below to show that IFC still achieves performance gains under different AoI traffic loads.

We vary the AoI traffic load by adjusting the AoI load level – the fraction of AoI traffic in the total traffic – and run the experiments under FIFO and standard IFC (SDM with $\gamma = 0.5$). The normalized results are shown in Table III, where we normalize the FIFO results by IFC. In most networks and load levels, IFC achieves both higher throughput and shorter AoI. In some cases, IFC trades a little throughput for significant reductions in AoI (2 to 9% throughput in exchange for 1.33 to $4.38\times$ shorter AoI). We remark that one can further adjust the AoI ratio γ in those cases for other trade-offs. Overall, IFC helps regardless of the AoI load level.

VIII. Related Work

Data freshness is a well-recognized requirement to achieve safety and stability in many communities, including connected vehicles/satellites [6], [15], [7], [8], [9], [10], real-time database [67], [68], [44], robotics [16], [17], [18], [19], and

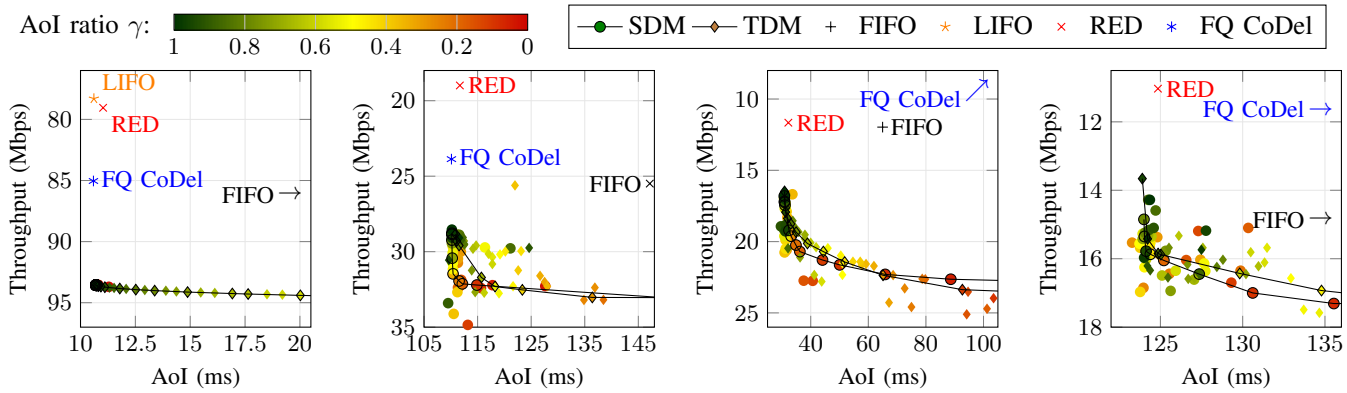

(a) Single link (b) SWAN (c) Internet2 (d) B4

Figure 15: We emulate SDM (round marks) and TDM (diamond marks) in Mininet and compare the methods with state-of-the-art queueing disciplines, including FIFO, RED, and FQ-CoDel. We also verify using the single-link setting that LIFO could be beneficial for AoI but significantly harmful for throughput. Notice that the horizontal axis is reversed, and the closer to the bottom-left corner implies better performance. We also draw the trade-off curves of SDM and TDM. As shown in the figures, SDM and TDM can achieve higher throughput with lower AoI using an appropriately chosen AoI ratio. As expected, the trade-off curves show that SDM and TDM strive for shorter AoI under larger AoI ratio γ.

Table II: The normalized throughput (Thr) and AoI of IFC (with SDM as the scheduler) combined with existing methods (mean $\pm$ standard deviation, normalized by standard IFC). The results show that IFC improves the throughput by 10% to 40% and reduces the AoI up to 50%.

(a) Internet2

Method	without IFC		with IFC(SDM)	
	Thr	AoI	Thr	AoI
FIFO	0.72 ± 0.05	2.14 ± 0.13	1.00 ± 0.00	1.00 ± 0.00
FQ CoDel	0.52 ± 0.08	889.42 ± 157.41	0.67 ± 0.08	686.96 ± 159.86
RED	0.65 ± 0.04	1.03 ± 0.02	0.81 ± 0.05	0.99 ± 0.02
BBR	0.72 ± 0.06	2.14 ± 0.13	1.01 ± 0.06	1.00 ± 0.02
DCTCP	0.74 ± 0.05	1.39 ± 0.05	1.03 ± 0.06	1.00 ± 0.02

(b) B4

Method	without IFC		with IFC(SDM)	
	Thr	AoI	Thr	AoI
FIFO	0.92 ± 0.07	1.21 ± 0.04	1.00 ± 0.00	1.00 ± 0.00
FQ CoDel	0.71 ± 0.08	37.76 ± 19.41	0.87 ± 0.09	23.38 ± 9.69
RED	0.70 ± 0.05	1.01 ± 0.02	0.78 ± 0.06	1.00 ± 0.02
BBR	0.92 ± 0.11	1.19 ± 0.13	0.99 ± 0.12	0.99 ± 0.10
DCTCP	0.89 ± 0.07	1.03 ± 0.02	0.99 ± 0.08	1.00 ± 0.02

Table III: Normalized throughput and AoI under different load level of AoI traffic (normalized by standard IFC). For all cases, IFC improves both throughput and AoI, or trades a little throughput to shorten the AoI significantly.

(a) Normalized throughput (the smaller the better).

Network	AoI Load Level				
	10%	1%	0.1%	0.01%	0.001%
Single Link	0.93	1.00	1.00	1.00	0.99
SWAN	0.86	1.04	1.05	0.99	1.02
Internet2	0.84	1.09	0.99	0.99	1.00
B4	0.99	1.00	1.00	0.90	0.93

(b) Normalized AoI (the larger the better).

Network	AoI Load Level				
	10%	1%	0.1%	0.01%	0.001%
Single Link	13.14	56.79	33.70	7.38	1.71
SWAN	1.36	2.83	4.05	2.87	1.33
Internet2	2.20	4.38	5.63	2.50	1.22
B4	1.23	1.96	2.42	1.93	1.21

power systems [11], [46], [12], [47], [13], [14]. State-of-the-art designs use networks for connectivity only and strive for freshness at the end systems by active probing [43], [44], [45], dedicated links [46], [47], or local decision [9].

The goal of improving freshness of flows in a network has also captured the interest of the information theory community over the past decade since the initial work of [24], [25] introduced the concept of the AoI. So far, most research papers in this area focus on theoretical analysis and consider the updates generated by a stochastic source going through different queueing systems, including one-queue systems ($M/M/1$, $M/D/1$, $D/M/1$ [25], $M/M/1/1$, and $M/M/1/2*$ [29]) and multihop networks [33]. This literature has provided many insights for design. For example, when multiple sources are present, [31] proposes to keep only the freshest update in the queue and simulation results demonstrate that the approach can effectively reduce the queueing delay. Another example is that the zero-wait policy, a.k.a., the work-conservation policy, is usually not AoI optimal [69], [32]. Instead, a "lazy" policy performs better [69]. We refer the reader to [26] for a more comprehensive survey on the prior work on AoI.

One recent set of papers that is related to ours focuses on analytic studies of the interplay between freshness and traditional measures, like throughput and energy. For example, [70], [71], [72] focus on scheduling and queueing disciplines for flows that care about a combination of freshness and throughput. [73], [74] study the energy cost of maintaining freshness. We remark that the problems investigated in these papers are different from ours in that they impose constraints to maintain throughput/energy level for the AoI flows. In contrast, in our work, the AoI flows are interested in AoI only, and a separate class of flows is interested in throughput. We

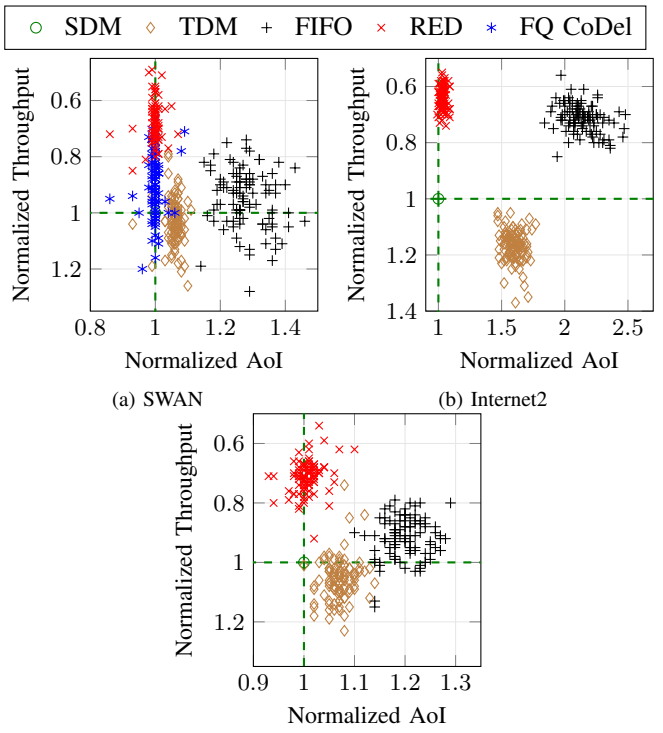

Figure 16: We run 100 randomly generated flow patterns and normalize the results by the standard IFC (SDM with $\gamma = 0.5$). FQ-CoDel achieves much lower throughput and higher AoI in both Internet2 and B4 that exceed the top-right corners of the figures.

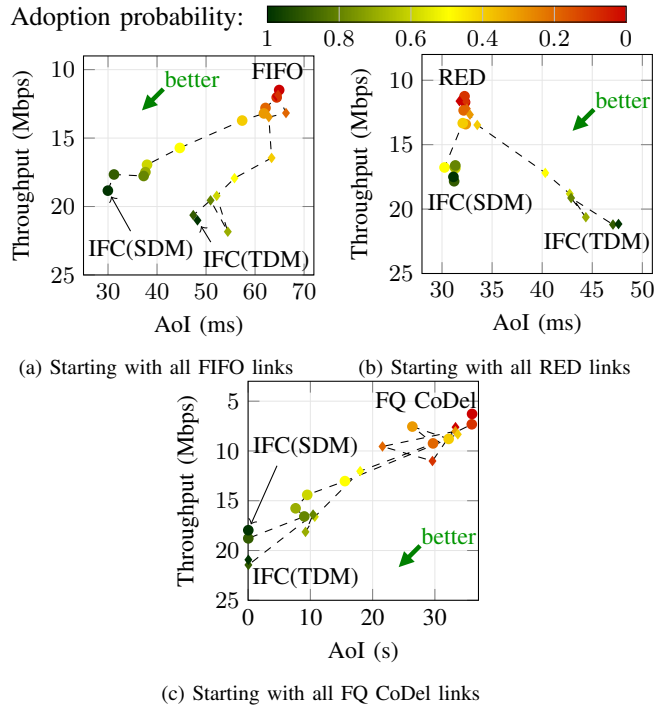

Figure 17: Random links adopting IFC according to the adoption probability. Originally, all the links obey some queueing discipline (FIFO, RED, or FQ CoDel). As more links adopt IFC, the performance improves.

emphasize that, to the best of our knowledge, our paper is the first contribution to this literature that includes implementation experiments.

Regarding implementation, our approach to freshness is related to two ideas: active queue management and heterogeneous flow scheduling. Active queue management methods control the queue size through packet drop or explicit congestion notification. They aim to avoid congestion [64], [75] and bufferbloat [76], [77], [65]. Although those methods also drop packets, they do not consider the heterogeneity of flows, neither do they differentiate packets from the same flow according to their properties such as freshness.

On the other hand, heterogeneous flow scheduling systems acknowledge the variety of objectives among flows and exploit their differences to get better performance. For instance, HULL [78] deals with throughput and ultra-low latency applications. pFabric [52] and PIAS [53] reduce flow completion time by differentiating short and long flows. D^3 [54] satisfies deadlines of alive flows by quenching the flows past due. NetStitcher [79] improves average utilization by manipulating bulk traffic. All of the systems above leverage the diversity of the flows to achieve their design goals, but our IFC is the first to address the trade-off between throughput and freshness.

IX. Concluding Remarks

In summary, our proposed In-network Freshness Control (IFC) presents the first implementation of a design for improving freshness of network flows while maintaining high throughput for legacy traffic via a fine-grained control of the trade-off between the two quantities. IFC alleviates the requirements on the endhosts, and it addresses the network sharing problem naturally that is not easily handleable for endhost-based freshness control. We implement our design as Linux kernel modules and show through emulations that IFC can reduce age of information (AoI) dramatically, while simultaneously improving throughput (in many cases), or costing throughput by only a small amount (in the worst case). Further, we demonstrate that IFC can incorporate existing methods for legacy traffic and does not require full-scale adoption to realize performance benefits.

There are a number of interesting issues that remain to study in followup work. Firstly, by dropping packets, IFC no longer guarantees reliability for AoI flows, which can raise security concerns. Currently, there are two ways to handle such concerns in the IFC design – either marking security-emphasized packets as LDA or setting the AoI ratio to 1 to fully prioritize AoI flows. Though these approaches would work, more sophisticated measures to balance freshness and security are important to explore. Secondly, we give analytic guidelines for setting the AoI ratio in IFC based on modeling the network as a single link. Of course, real networks can be more complicated, and load balancing under different LDA/AoI traffic composition is itself an interesting research direction. Also, in more general situations, one may wonder if setting the AoI ratio differently at each port would give a better trade-off curve. Thirdly, all the designs we consider are work-conserving, and while this is true with traditional measures like

throughput and delay, optimal AoI designs may not be work-conserving. As we have seen with the pulse train in Section V, it may be worth waiting instead of starting transmission even if the AoI queue is nonempty and the link is idle. It will be interesting to investigate designs that are not work-conserving in future work in order to see what improvements are possible. Finally, although our IFC design can be implemented in both DPDK and Linux kernel, for even higher-performance networks, people are interested in deploying solutions to the Protocol Independent Switch Architecture (PISA). Since IFIL may replace enqueued packets and PISA in general does not touch enqueued packets, deploying IFC to PISA requires more research.

REFERENCES

[1] Google Stadia. Cloud gaming platform. [Online]. Available: https://store.google.com/magazine/stadia

[2] Facebook live. [Online]. Available: https://live.fb.com

[3] YouTube TV. [Online]. Available: https://tv.youtube.com

[4] Microsoft Azure IoT. [Online]. Available: https://azure.microsoft.com/en-us/overview/iot

[5] K. Ashton *et al.*, "That 'Internet of Things' thing," *RFID journal*, vol. 22, no. 7, pp. 97–114, 2009.

[6] T. Gandhi and M. M. Trivedi, "Pedestrian protection systems: Issues, survey, and challenges," *IEEE Transactions on Intelligent Transportation Systems*, vol. 8, no. 3, pp. 413–430, 2007.

[7] R. Talak, S. Karaman, and E. Modiano, "Speed limits in autonomous vehicular networks due to communication constraints," in *Proc. IEEE CDC*. IEEE, 2016, pp. 4998–5003.

[8] J.-D. M. M. Biomo, T. Kunz, and M. St-Hilaire, "Routing in unmanned aerial ad hoc networks: Introducing a route reliability criterion," in *2014 7th IFIP wireless and mobile networking conference (WMNC)*. IEEE, 2014, pp. 1–7.

[9] C.-J. Hoel, K. Driggs-Campbell, K. Wolff, L. Laine, and M. Kochenderfer, "Combining planning and deep reinforcement learning in tactical decision making for autonomous driving," *IEEE Transactions on Intelligent Vehicles*, vol. 9, no. 3, pp. 118–122, 2019.

[10] R. Talak, S. Karaman, and E. Modiano, "Optimizing information freshness in wireless networks under general interference constraints," *IEEE/ACM Trans. Netw.*, 2019.

[11] A. G. Phadke, "Synchronized phasor measurements in power systems," *IEEE Computer Applications in power*, vol. 6, no. 2, pp. 10–15, 1993.

[12] V. Terzija, G. Valverde, D. Cai, P. Regulski, V. Madani, J. Fitch, S. Skok, M. M. Begovic, and A. Phadke, "Wide-area monitoring, protection, and control of future electric power networks," *Proceedings of the IEEE*, vol. 99, no. 1, pp. 80–93, 2011.

[13] G. Celli, P. A. Pegoraro, F. Pilo, G. Pisano, and S. Sulis, "DMS cyber-physical simulation for assessing the impact of state estimation and communication media in smart grid operation," *IEEE Transactions on Power Systems*, vol. 29, no. 5, pp. 2436–2446, 2014.

[14] A. Primadianto and C.-N. Lu, "A review on distribution system state estimation," *IEEE Transactions on Power Systems*, vol. 32, no. 5, pp. 3875–3883, 2016.

[15] S. Bandyopadhyay, R. Foust, G. Subramanian, S.-J. Chung, and F. Hadaegh, "Review of formation flying and constellation missions using nanosatellites," *Journal of Spacecraft and Rockets*, vol. 53, pp. 1–12, Mar 2016.

[16] A. R. Lawton, B. J. Young, and R. W. Beard, "A decentralized approach to elementary formation manoeuvres," in *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings*, vol. 3, Apr 2000, pp. 2728–2733.

[17] S. Bandyopadhyay, S. Chung, and F. Y. Hadaegh, "Probabilistic and distributed control of a large-scale swarm of autonomous agents," *IEEE Transactions on Robotics*, vol. 33, no. 5, pp. 1103–1123, Oct 2017.

[18] D. Tardioli, D. Sicignano, L. Riazuelo, A. Romeo, J. L. Villarroel, and L. Montano, "Robot teams for intervention in confined and structured environments," *J. Field Robotics*, vol. 33, pp. 765–801, 2016.

[19] "DARPA Subterranean (SubT) challenge," https://www.subtchallenge.com/, accessed: 2020-02-20.

[20] R. M. Allen and H. Kanamori, "The potential for earthquake early warning in southern california," *Science*, vol. 300, no. 5620, pp. 786–789, 2003.

[21] O. Kamigaichi, "JMA earthquake early warning," *Journal of Japan Association for Earthquake Engineering*, vol. 4, no. 3, pp. 134–137, 2004.

[22] N.-C. Hsiao, Y.-M. Wu, T.-C. Shin, L. Zhao, and T.-L. Teng, "Development of earthquake early warning system in taiwan," *Geophysical research letters*, vol. 36, no. 5, pp. 91–100, 2009.

[23] Cisco Visual Networking Index, "Global mobile data traffic forecast update, 2017 – 2022," *Cisco White Paper*, 2019.

[24] S. Kaul, M. Gruteser, V. Rai, and J. Kenney, "Minimizing age of information in vehicular networks," in *Proc. IEEE SECON*. IEEE, 2011, pp. 350–358.

[25] S. Kaul, R. Yates, and M. Gruteser, "Real-time status: How often should one update?" in *Proc. IEEE INFOCOM*. IEEE, 2012, pp. 2731–2735.

[26] R. D. Yates and S. K. Kaul, "The age of information: Real-time status updating by multiple sources," *IEEE Trans. Inf. Theory*, vol. 65, no. 3, pp. 1807–1827, 2019.

[27] T. Shreedhar, S. K. Kaul, and R. D. Yates, "ACP: Age control protocol for minimizing age of information over the Internet," in *Proc. IEEE MobiCom*, 2018, pp. 699–701.

[28] I. Kadota, M. S. Rahman, and E. Modiano, "Age of information in wireless networks: From theory to implementation," in *Proc. IEEE MobiCom*, 2020, pp. 1–3.

[29] M. Costa, M. Codreanu, and A. Ephremides, "Age of information with packet management," in *Proc. IEEE ISIT*. IEEE, 2014, pp. 1583–1587.

[30] L. Huang and E. Modiano, "Optimizing age-of-information in a multi-class queueing system," in *Proc. IEEE ISIT*. IEEE, 2015, pp. 1681–1685.

[31] N. Pappas, J. Gunnarsson, L. Kratz, M. Kountouris, and V. Angelakis, "Age of information of multiple sources with queue management," in *Proc. IEEE ICC*. IEEE, 2015, pp. 5935–5940.

[32] Y. Sun, E. Uysal-Biyikoglu, R. D. Yates, C. E. Koksal, and N. B. Shroff, "Update or wait: How to keep your data fresh," *IEEE Trans. Inf. Theory*, vol. 63, no. 11, pp. 7492–7508, 2017.

[33] A. M. Bedewy, Y. Sun, and N. B. Shroff, "Age-optimal information updates in multihop networks," in *Proc. IEEE ISIT*. IEEE, 2017, pp. 576–580.

[34] I. Kadota, A. Sinha, E. Uysal-Biyikoglu, R. Singh, and E. Modiano, "Scheduling policies for minimizing age of information in broadcast wireless networks," *IEEE/ACM Trans. Netw.*, vol. 26, no. 6, pp. 2637–2650, 2018.

[35] V. Kavitha, E. Altman, and I. Saha, "Controlling packet drops to improve freshness of information," *arXiv preprint arXiv:1807.09325*, 2018.

[36] M. Liu, L. Luo, J. Nelson, L. Ceze, A. Krishnamurthy, and K. Atreya, "IncBricks: Toward in-network computation with an in-network cache," in *Proc. ASPLOS*, 2017, pp. 795–809.

[37] X. Jin, X. Li, H. Zhang, R. Soulé, J. Lee, N. Foster, C. Kim, and I. Stoica, "NetCache: Balancing key-value stores with fast in-network caching," in *Proc. ACM SOSP*, 2017, pp. 121–136.

[38] A. Sapio, I. Abdelaziz, A. Aldilaijan, M. Canini, and P. Kalnis, "In-network computation is a dumb idea whose time has come," in *Proc. HotNets*. ACM, 2017, pp. 150–156.

[39] N. Cardwell, Y. Cheng, C. S. Gunn, S. H. Yeganeh, and V. Jacobson, "BBR: Congestion-based congestion control," *Queue*, vol. 14, no. 5, p. 50, 2016.

[40] M. Alizadeh, A. Greenberg, D. A. Maltz, J. Padhye, P. Patel, B. Prabhakar, S. Sengupta, and M. Sridharan, "Data center TCP (DCTCP)," *ACM SIGCOMM CCR*, vol. 41, no. 4, pp. 63–74, 2011.

[41] J. Sherry, S. Hasan, C. Scott, A. Krishnamurthy, S. Ratnasamy, and V. Sekar, "Making middleboxes someone else's problem: Network processing as a cloud service," *ACM SIGCOMM CCR*, vol. 42, no. 4, pp. 13–24, 2012.

[42] Z. A. Qazi, C.-C. Tu, L. Chiang, R. Miao, V. Sekar, and M. Yu, "SIMPLE-fying middlebox policy enforcement using SDN," *ACM SIGCOMM CCR*, vol. 43, no. 4, pp. 27–38, 2013.

[43] S. Hu, S. Yao, H. Jin, Y. Zhao, Y. Hu, X. Liu, N. Naghibolhosseini, S. Li, A. Kapoor, W. Dron *et al.*, "Data acquisition for real-time decision-making under freshness constraints," in *2015 IEEE Real-Time Systems Symposium*. IEEE, 2015, pp. 185–194.

[44] J.-E. Kim, T. Abdelzaher, L. Sha, A. Bar-Noy, and R. Hobbs, "Sporadic decision-centric data scheduling with normally-off sensors," in *2016 IEEE Real-Time Systems Symposium (RTSS)*. IEEE, 2016, pp. 135–145.

[45] A. Dhekne, A. Chakraborty, K. Sundaresan, and S. Rangarajan, "TrackIO: Tracking first responders inside-out," in *Proc. USENIX NSDI*, 2019, pp. 751–764.

[46] B. Naduvathuparambil, M. C. Valenti, and A. Feliachi, "Communication delays in wide area measurement systems," in *Proceedings of the Thirty-*

Fourth Southeastern Symposium on System Theory (Cat. No. 02EX540). IEEE, 2002, pp. 118–122.

[47] D. Ghosh, T. Ghose, and D. K. Mohanta, "Communication feasibility analysis for smart grid with phasor measurement units," *IEEE Transactions on Industrial Informatics*, vol. 9, no. 3, pp. 1486–1496, 2013.

[48] M. Handley, C. Raiciu, A. Agache, A. Voinescu, A. W. Moore, G. Antichi, and M. Wójcik, "Re-architecting datacenter networks and stacks for low latency and high performance," in *Proc. ACM SIGCOMM*. ACM, 2017, pp. 29–42.

[49] A. Saeed, Y. Zhao, N. Dukkipati, M. Ammar, E. Zegura, K. Harras, and A. Vahdat, "Eiffel: Efficient and flexible software packet scheduling," in *Proc. USENIX NSDI*. USENIX Association, 2019, pp. 17–31.

[50] V. Shrivastav, "Fast, scalable, and programmable packet scheduler in hardware," in *Proc. ACM SIGCOMM*. ACM, 2019, pp. 367–379.

[51] N. Dukkipati and N. McKeown, "Why flow-completion time is the right metric for congestion control," *ACM SIGCOMM CCR*, vol. 36, no. 1, pp. 59–62, 2006.

[52] M. Alizadeh, S. Yang, M. Sharif, S. Katti, N. McKeown, B. Prabhakar, and S. Shenker, "pFabric: Minimal near-optimal datacenter transport," *ACM SIGCOMM CCR*, vol. 43, no. 4, pp. 435–446, 2013.

[53] W. Bai, K. Chen, H. Wang, L. Chen, D. Han, and C. Tian, "Information-agnostic flow scheduling for commodity data centers," in *Proc. USENIX NSDI*, 2015, pp. 455–468.

[54] C. Wilson, H. Ballani, T. Karagiannis, and A. Rowtron, "Better never than late: Meeting deadlines in datacenter networks," *ACM SIGCOMM CCR*, vol. 41, no. 4, pp. 50–61, 2011.

[55] B. Vamanan, J. Hasan, and T. Vijaykumar, "Deadline-aware datacenter TCP (D2TCP)," *ACM SIGCOMM CCR*, vol. 42, no. 4, pp. 115–126, 2012.

[56] A. Gushchin, S.-H. Tseng, and A. Tang, "Optimization-based network flow deadline scheduling," in *Proc. IEEE ICNP*, nov 2016.

[57] A. Franco, E. Fitzgerald, B. Landfeldt, N. Pappas, and V. Angelakis, "LUPMAC: A cross-layer MAC technique to improve the age of information over dense WLANs," in *Proc. IEEE ICT*, 2016, pp. 1–6.

[58] Data plane development kit. [Online]. Available: http://dpdk.org

[59] Mininet. [Online]. Available: http://mininet.org/

[60] C.-Y. Hong, S. Kandula, R. Mahajan, M. Zhang, V. Gill, M. Nanduri, and R. Wattenhofer, "Achieving high utilization with software-driven WAN," *ACM SIGCOMM CCR*, vol. 43, no. 4, pp. 15–26, 2013.

[61] The Internet2 network. [Online]. Available: http://www.internet2.edu/

[62] S. Jain, A. Kumar, S. Mandal, J. Ong, L. Poutievski, A. Singh, S. Venkata, J. Wanderer, J. Zhou, M. Zhu, J. Zolla, U. Hölzle, S. Stuart, and A. Vahdat, "B4: Experience with a globally-deployed software defined WAN," *ACM SIGCOMM CCR*, vol. 43, no. 4, pp. 3–14, 2013.

[63] C.-Y. Hong, S. Mandal, M. Al-Fares, M. Zhu, R. Alimi, K. Naidu B., C. Bhagat, S. Jain, J. Kaimal, S. Liang, K. Mendelev, S. Padgett, F. Rabe, S. Ray, M. Tewari, M. Tierney, M. Zahn, J. Zolla, J. Ong, and A. Vahdat, "B4 and after: Managing hierarchy, partitioning, and asymmetry for availability and scale in Google's software-defined WAN," in *Proc. ACM SIGCOMM*, 2018.

[64] V. Jacobson, B. Braden, S. Floyd, B. Davie, D. Estrin, J. Crowcroft, and S. Deering, "RFC 2309: Recommendations on queue management and congestion avoidance in the Internet," 1998.

[65] D. Taht, J. Gettys, E. Dumazet, T. Hoeiland-Joergensen, E. Dumazet, P. McKenney, J. Gettys, T. Hoeiland-Joergensen, and P. McKenney, "RFC 8290: The flow queue codel packet scheduler and active queue management algorithm," 2018.

[66] The CAIDA UCSD anonymized Internet traces 2018 and 2019. [Online]. Available: http://www.caida.org/data/passive/passive_dataset.xml

[67] B. Adelberg, H. Garcia-Molina, and B. Kao, "Applying update streams in a soft real-time database system," in *Proc. ACM SIGMOD*, vol. 24, no. 2. ACM, 1995, pp. 245–256.

[68] K.-D. Kang, S. H. Son, and J. A. Stankovic, "Managing deadline miss ratio and sensor data freshness in real-time databases," *IEEE Transactions on Knowledge and Data Engineering*, vol. 16, no. 10, pp. 1200–1216, 2004.

[69] R. D. Yates, "Lazy is timely: Status updates by an energy harvesting source," in *Proc. IEEE ISIT*. IEEE, 2015, pp. 3008–3012.

[70] A. M. Bedewy, Y. Sun, and N. B. Shroff, "Optimizing data freshness, throughput, and delay in multi-server information-update systems," in *Proc. IEEE ISIT*. IEEE, 2016, pp. 2569–2573.

[71] I. Kadota, A. Sinha, and E. Modiano, "Optimizing age of information in wireless networks with throughput constraints," in *Proc. IEEE INFOCOM*. IEEE, 2018, pp. 1844–1852.

[72] ——, "Scheduling algorithms for optimizing age of information in wireless networks with throughput constraints," *IEEE/ACM Trans. Netw.*, vol. 27, no. 4, pp. 1359–1372, 2019.

[73] E. Altman, R. El-Azouzi, D. S. Menasche, and Y. Xu, "Forever young: Aging control for hybrid networks," in *Proceedings of the Twentieth ACM International Symposium on Mobile Ad Hoc Networking and Computing*, 2019, pp. 91–100.

[74] E. T. Ceran, D. Gündüz, and A. György, "Average age of information with hybrid ARQ under a resource constraint," *IEEE Transactions on Wireless Communications*, vol. 18, no. 3, pp. 1900–1913, 2019.

[75] W.-C. Feng, D. D. Kandlur, D. Saha, and K. G. Shin, "A self-configuring RED gateway," in *Proc. IEEE INFOCOM*. IEEE, 1999, pp. 1320–1328.

[76] K. Nichols and V. Jacobson, "Controlling queue delay," *Communications of the ACM*, vol. 55, no. 7, pp. 42–50, 2012.

[77] R. Pan, P. Natarajan, C. Piglione, M. S. Prabhu, V. Subramanian, F. Baker, and B. VerSteeg, "PIE: A lightweight control scheme to address the bufferbloat problem," in *Proc. IEEE HPSR*, 2013, pp. 148–155.

[78] M. Alizadeh, A. Kabbani, T. Edsall, B. Prabhakar, A. Vahdat, and M. Yasuda, "Less is more: Trading a little bandwidth for ultra-low latency in the data center," in *Proc. USENIX NSDI*, 2012.

[79] N. Laoutaris, M. Sirivianos, X. Yang, and P. Rodriguez, "Inter-datacenter bulk transfers with NetStitcher," *ACM SIGCOMM CCR*, vol. 41, no. 4, pp. 74–85, 2011.