arXiv:2610.07230v1 [cs.NI] 5 Oct 2026

Ofan: Optimal Load Balancing for AI Training

Sarah McClure[1] Evyatar Cohen[2] Jakob Krebs[2] Alexander Shpiner[3]
Mark Silberstein[2,3] Scott Shenker[1,4] Sylvia Ratnasamy[1] Isaac Keslassy[1,2]

[1] *UC Berkeley* [2] *Technion* [3] *NVIDIA* [4] *ICSI*

Abstract

The extreme collective completion time (CCT) demands of AI workloads challenge existing packet spraying algorithms, which can have trouble efficiently load-balancing workloads that are sent at full line rates.

We trace this to a structural cause: on a fat tree, once a packet picks its upward path, the downward path to its destination is unique, so destination-oblivious schemes cannot undo the imbalance it creates. We prove that such schemes can suffer from $\Theta(\sqrt{m})$ queueing for messages of size m, thus eventually triggering rate reductions by the congestion control. Instead, we suggest OFAN, a switch-based destination-aware LB scheme that can reach $O(1)$ queueing. We also present its POFAN variant that fits the pipe-based architecture of current switches. Our P4 implementation shows that it consumes modest resources. An end-to-end FSDP2 evaluation with Llama-3 405B-parameter models shows that OFAN cuts CCT inflation by 16–39× when compared to existing algorithms.

1 Introduction

Thirty years of load balancing. Datacenter networks offer a rich set of parallel paths, and a long line of load-balancing (LB) algorithms has been proposed to exploit them. That line was long constrained by the transport protocol's intolerance for reordering, and so operated at a *coarse* granularity: per-flow hashing [36] and its weighted variants [60, 88], subflows [1, 71, 74], host-based path adaptation [31, 72, 77], and flowlet switching [41, 42, 82]. All of these pin a unit of traffic, whether a flow, a subflow, or a flowlet, to a single path for its lifetime. On the contrary, relying on NICs that accept out-of-order delivery [18, 62, 78], *spraying* algorithms load-balance at the granularity of a single packet. These include (i) random packet spraying (RPS) [26] as in Alibaba's Stellar [56], (ii) host-based congestion-aware spraying algorithms like REPS and MRC [5, 8, 49] as in the Ultra Ethernet Consortium [18], and (iii) switch-based algorithms that steer packets based on local congestion like DRILL [33] and Nvidia's Spectrum-X [45, 63, 64] adaptive routing (AR) [27, 34], or also take into account congestion at neighboring switches [3, 70], like Broadcom's GLB and ByteDance's SGLB.

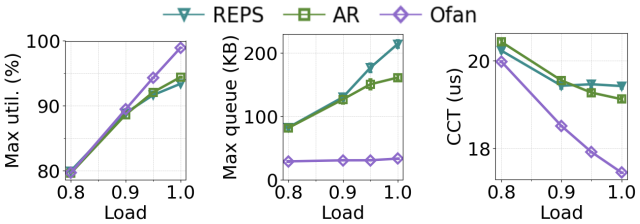

(a) Max utilization **(b) Max queue size** **(c) CCT**

Figure 1: Utilization and max queue size as load increases.

AI training. Recent AI training workloads have new demands. They focus on minimizing collective completion time (CCT), i.e., the worst-case flow completion time among all flows in a collective, and therefore push every source to send at 100% of its line rate [21, 22, 40].

Unfortunately, as Fig. 1 shows, existing algorithms fail to meet this demand. We simulate a three-level fat-tree with a random permutation workload using 1 MB messages (as detailed in §6), and sweep the offered load toward 100%. We also remove congestion control entirely, pacing each source at a fixed rate dictated by the load, so that we only measure the effects of load balancing (§2.2). The results display the limits of current state-of-the-art spraying algorithms. Fig. 1(a) shows that for REPS and AR, utilization cannot keep up with load beyond load = 0.9 (it disconnects from the $y = x$ diagonal), and at load 1 it saturates below 95%. Likewise, Fig. 1(b) illustrates how their queue sizes of REPS and AR grow significantly despite the small message size. Finally, Fig. 1(c) shows that their CCTs improve less beyond a load of 0.9, also pointing to some saturation. In this paper, we claim that all these issues come from a sub-optimal load balancing: as packets do not reach all queues equally, some queues become overloaded (high queue sizes) while others remain underutilized (low utilization), thus impacting the CCT. We introduce a new LB algorithm, OFAN. These plots show how its efficient LB

properties enable OFAN to achieve about 99% utilization, much smaller queue sizes, and a clearly better CCT. Note that once congestion control is switched back on, we would expect ECN congestion signals (and PFC pauses for larger messages) to occur for REPS and AR, yielding an even worse CCT gap.

OFAN. In the paper, we introduce OFAN, a new LB algorithm that relies on three key ideas:

(1) Per-destination LB. In fat trees, LB has freedom only on the way up, while the path down to a destination is unique. Existing switch-based LB schemes typically select outgoing port to optimize upward LB. This balances the northbound links beautifully. But once a packet has picked its upward path, no downstream switch can correct a bad choice. A locally-efficient but destination-oblivious LB scheme therefore spreads each destination's traffic unevenly when going up, and that unevenness cannot be addressed when going down.

In this paper, we consider random permutation workloads with message size m and assume that all hosts blast traffic at 100% load. We prove that randomized spraying algorithms incur an average queue size of $\Omega(\sqrt{m})$ (Thm. 1). We also show that in evaluations, additional destination-oblivious algorithms follow a similar heuristic growth. In practice, these bounds mean that as message sizes grow, queues grow, and therefore congestion control algorithms (CCAs) will be forced by ECNs and PFCs to slow down sending rates and increase CCTs, hurting AI training performance.

Instead, inspired by host-based DRB [16] and ETHEREAL [1], we propose a *destination-based rotation* (DR) discipline, in which consecutive packets to a given destination are deterministically rotated across paths. Using $ND/D/1$ queueing models [58, 76], we show that destination-based rotation can achieve optimal $O(1)$ queue bounds even at 100% loads (Thm. 3).

(2) Switch-based LB. We argue that the switches are a better home for DR than the hosts. First, the switch already knows the full topology and its link failures using BGP, while the host must either implement BGP in the NIC or keep re-inferring it indirectly. Second, and this is the property we find most surprising, a switch can *consolidate* traffic for a given destination, making packets from different sources share the same rotation pointer. A host, seeing only its own flows, cannot do this.

(3) Pipe-based LB. Finally, we show that we can exploit current *pipe*-based switch implementations (called pipes in Broadcom's Tomahawk switches and slices in Cisco's Silicon One switches). Pipes aggregate a few ports (e.g., 4 ports at 800Gb/s each) and guarantee atomic Read-Modify-Write (RMW) accesses to local registers at line rate. Thus, we introduce POFAN, a more practical variant of OFAN that stores a local pointer at each pipe for each destination instead of a shared pointer accessible across the whole switch.

Contributions. In summary, we make the following contributions:

- We demonstrate that existing randomized-spraying algorithms suffer from an intrinsic queuing lower bound that grows with the message size m as $\Omega(\sqrt{m})$. We also demonstrate that some LB algorithms like join-the-shortest-queue (JSQ) even suffer from newly discovered synchronization issues at full line rate, and suffer from a $\Theta(m)$ queue growth. On the contrary, we show that destination-based rotation achieves $O(1)$ queue bounds that enable larger messages without queue growth (§3).
- We present OFAN, a switch-based DR, and explain the intuition behind its consolidation properties (§4).
- We further present how to implement its variant POFAN in a pipe-based switch, and how it requires limited resources in a P4 implementation for an Intel Tofino switch (§5).
- Finally, we show how network size, message size, packet size, collective type, link failures, topology, incremental deployment and congestion control affect the results. We find that OFAN and POFAN keep outperforming existing algorithms across a wide variety of conditions. We further evaluate an AI training scenario with Llama-3 (405B parameters) using FSDP2, and rely on two different recent CCAs designed for spraying: UEC's NSCC [18, 81] and MSwift [32]. OFAN reduces the CCT increase beyond a CCT lower-bound by at least 16× with NSCC and 39× with MSwift when compared to AR and REPS (and POFAN with four ports per pipe reduces by 15× and 36×, respectively). (§6).

The rest of the paper is organized as follows: §2 states the workload properties our analysis assumes and the methodology that isolates the load-balancing layer; §3 explain the intuition behind per-destination LB, proves the limits of most existing algorithms, and offers DR as a way forward; §4 presents OFAN and §5 its realization in switch hardware; §6 evaluates it, and §7 discusses what remains open.[1]

2 Background and Methodology

2.1 AI Training Workloads

AI training alternates rounds of computation with rounds of communication, where the communication is one of a few standard collectives, such as AllReduce, AllGather, and AllToAll, drawn from a small set of libraries [21, 22, 52]. Two properties of this traffic matter for load balancing:

Uniform. The nodes in a collective run the same communication library and the same transport and NICs, and most collectives are symmetric: every participant has the same amount of data to send, with the same packet size.

[1] This paper is a follow-up on [57], which had two parts: (1) systematically evaluating leading load-balancing designs, while decoupling them from specific congestion control and loss recovery stacks; and (2) introducing and studying OFAN. This paper focuses on the second part.

Synchronous. All nodes begin at roughly the same time [1]. This cuts both ways. It helps load balancing, since there are no newly arriving flows to adapt to mid-collective. But it also means that every source steps through its packets in lockstep, which can cause a load balancer's port choices to synchronize across sources and degrade catastrophically (as shown in §3).

Together with the last-byte nature of CCT, these properties also mean that the traffic is long-lived and gapless: techniques that rebalance during idle periods, such as PLB or flowlet switching, are less helpful. We also assume the cluster runs the job in isolation, as is common at scale [31, 47, 71].

Throughout this paper, we will focus on three types of generic traffic matrices that reflect generic AI training workloads [85]: (i) *random permutations*, (ii) *all-to-all*, and (iii) *MoE (all-to-all-v)*. In particular, we will show that all-to-all is the easier one to load-balance, and all-to-all-v suffers from a large incast that reduces the impact of load-balancing. Therefore, we will most focus on permutations, where the LB scheme can make the biggest impact.

2.2 Isolating the Load-Balancing Layer

Load balancing does not run alone: congestion control sets sending rates, and loss recovery decides what to retransmit. Both react to the queues that load balancing creates, so measuring a full stack conflates the three. Our methodology in the analysis is therefore to remove the other two and measure load balancing on its own (as in [57]). In evaluations (§6), we add CCAs back, and find that they do not change the main results, including the ranking among LB algorithms.

Congestion control. A perfect load balancer makes the network's entire capacity available, and an ideal congestion control algorithm (CCA) would provide each host with its fair share of that capacity. Under uniform AI collectives, the fair share is the same for all hosts. Thus, we cancel the CCA and allow instead each host to send at the maximum fraction of the line rate that the network can support, which we call $\rho_{\max}$. In a fault-free fabric, each host can blast its traffic at full line rate. Under failures, $\rho_{\max}$ is the highest rate at which all flows could send if traffic were balanced perfectly across all valid shortest paths. In addition, if a host has several flows, it simply performs round-robin (RR) to schedule them. Any queue that then forms, and any capacity left unused, is attributable to the load balancer and nothing else. This is what Fig. 1 measures.

Loss recovery. Similarly, we attempt to avoid any losses: we assume large buffers and do not consider failed links (unless mentioned).

This idealization is deliberate, and it is what lets us state an optimality criterion at all: a bound proved against a fixed-rate source is a property of the load balancer, not of a particular CCA's reaction to it. It is also what makes the criterion demanding, since there is no rate reduction available to rescue a scheme whose queues grow.

Table 1: Considered LB schemes.

Scheme	Where	Unit	Dest.-based	$q_{\uparrow}(m)$ (model)	$q_{\downarrow}(m)$ (model)	Assumptions
Destination-oblivious LB						
SIMPLE RR	switch	packet	✗	$\Theta(1)$	$\Theta(m)$	—
JSQ [27]	switch	packet	✗	$\Theta(1)$	$\Theta(m)$	—
CONGA [3]	switch	packet	✗	—	—	Pkt-based
CHOICE [33]	switch	packet	✗	—	—	—
DRILL [33]	switch	packet	✗	—	—	Pkt-based
DLB [10]	switch	packet	✗	—	—	Local SGLB
SGLB [70]	switch	packet	✗	—	—	—
RPS [26]	host	packet	✗	$\Theta(\sqrt{m})$	$\Theta(\sqrt{m})$	—
RSQ	switch	packet	✗	$\Theta(\sqrt{m})$	$\Theta(\sqrt{m})$	—
REPS [8]	host	packet	✗	—	—	ECN thres.: 10%
AR [63]	switch	packet	✗	—	—	Quanta: 5-10-20%
Destination-based LB						
DRB [16]	host	packet	✓	$\Theta(1)$	$\Theta(1)$	—
ETHEREAL [1]	host	subflow	✓	$\Theta(1)$	$\Theta(1)$	—
4OFAN (§5)	switch	packet	✓	$\Theta(1)$	$\Theta(1)$	—
OFAN (§4)	switch	packet	✓	$\Theta(1)$	$\Theta(1)$	—

3 Per-Destination LB

In this section, we present theoretical analysis that gives us a more rigorous understanding of optimality in load balancing approaches, and leads us to later propose Ofan, a new and optimal approach, in §4.

3.1 LB Algorithms

Table 1 lists load-balancing (LB) schemes that we consider in this section. They are all packet-based spraying algorithms except for ETHEREAL, which relies on subflows, and can be either host- or switch-based. It also lists some assumptions we made in simulations: e.g., the initial CONGA fits flowlets, and we adapt it to packet spraying to fit our workloads. Likewise, Broadcom's DLB is not described in detail, but the SGLB paper [70] sketches it as an SGLB-like algorithm that only uses the 2 local types of information out of the 5 described ones, so we implement it in this way. While all schemes take into account the destination to decide at each switch whether to send packets up or down, the *destination-oblivious LB* schemes consider all packets that go up in the same way.

We also introduce a few fundamental LB algorithms in the table, as they help convey intuition: SIMPLE RR consists of a simple round-robin model in which the switch load-balances all northbound packets across a set of next-hop ports in round-robin order. In *Join-the-Shortest-Queue (JSQ)*, each switch selects the next-hop port with the shortest queue for a northbound packet, with random tie-breaking between equal-length queues. It is a non-quantized version of AR. In addition, in *Randomized Switch Queueing (*RSQ*)*, each switch sends each northbound packet to a random port, and is analytically equivalent to host-based RPS, which changes the packet entropy (e.g., source port) to modify the ECMP header-based hashed port decision at each switch.

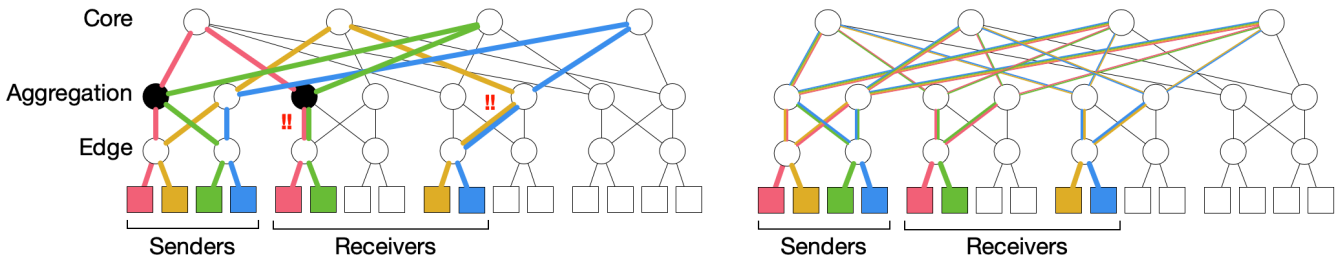

(a) Destination-oblivious LB (b) Destination-based LB

Figure 2: Impact of destination-based LB. The three-level fat tree with $k = 4$ includes 16 hosts, four pods (each with two edge and two aggregation switches), and four core switches. (a) Destination-oblivious LB can result in conflicts on the way down (red "!!"), while (b) destination-based LB avoids them.

3.2 The Downlink Problem

Throughout the paper, we consider three-level fat-tree topologies. In such topologies, LB flexibility exists strictly in upward paths, while downward paths to any destination are fixed. Consequently, while existing destination-oblivious LB algorithms may perfectly optimize for local LB properties on the way up, they may also distribute traffic bound for the same edge or aggregation switch unevenly across upward paths, thus creating downstream bottlenecks.

Fig. 2(a) illustrates how it is not sufficient for switches to simply balance traffic on the up routes. Each switch in the senders pod in the figure nicely sends half of its ingress traffic on each link, thus all up links are equally loaded. However, in this extreme example, all traffic to any destination edge switch ends up on the same aggregation switch in the senders pod. While this is fine for the up routes, this can cause collisions on the down routes, where destinations determine the route, not the LB. For example, the red and green paths meet at the first aggregation switch of the first pod, and therefore have to meet again at the first aggregation switch of the second pod (both drawn in black), and therefore conflict on the way down. In Fig. 2(b), however, switches not only balance load equally across all uplinks, they also ensure that the traffic sent to a particular destination is spread evenly. This ensures that there are no collisions on the down routes as an equal share of each flow traverses every possible path. Thus, this example shows how a destination-oblivious scheme can balance one side of the network and leave the other to chance, effectively causing statistical imbalance.

Fig. 3 quantifies the imbalance. We plot the worst-case link overload beyond the ideal, uniform amount, at each of the four inter-switch stages: edge to aggregation ($E \to A$), aggregation to core ($A \to C$), core to aggregation ($C \to A$), and aggregation to edge ($A \to E$). We assume a random inter-pod permutation that forces all flows to cross pods and in which all links are ideally equally loaded. The asymmetry is stark. For example, SIMPLE RR and JSQ achieve essentially perfect balance on the two *uplink* stages and then collapse on the two *downlink* stages; CHOICE and DRILL also exhibit similar profiles; the other destination-oblivious schemes suffer at both

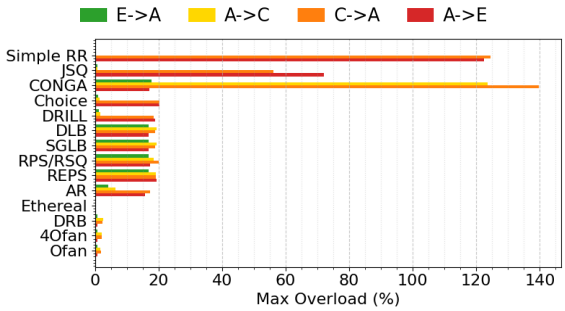

Figure 3: Worst-case link overload at each stage.

uplink and downlink. Only the destination-aware schemes are near-flat across all four, as we will detail below. In particular, ETHEREAL achieves ideal balance because it can decompose a single message into many subflows with one small packet in each.

3.3 Queue-Size Lower Bounds

We now want to provide a fundamental analysis of the performance of each LB scheme. We consider n hosts in a three-level k-ary fat tree without link failures. The hosts perform a random permutation collective, such that each host sends a single flow to a distinct random host (random derangement). All hosts send traffic at line rate, buffers are infinite, and there are no packet losses nor ACKs.

We define optimality in terms of the average queue length that a scheme produces. Collectives can have a large range of message sizes [85], and we want to ensure that as they get larger, the queue size does not grow without bound. For a message size m, let $q(m)$ be the average queue size across *all* switch queues, over the time it takes for *all* hosts to finish sending. An *optimal* load-balancing algorithm is one that achieves $q(m) = O(1)$, independent of m. We also decompose $q(m) = q_{\uparrow}(m) + q_{\downarrow}(m)$ to obtain the contribution of the uplink and downlink queues, respectively.

Queue size is the right target for several reasons. First, it is strongly correlated with CCT performance: given fixed propagation delay and our assumptions of uniform traffic, ideal

sending rates, and zero loss, queueing is the dominant variable component of the CCT, as we validate in §3.5. Second, it is also a dominant factor in congestion control, as it incurs ECNs, PFCs and buffer overflows. Third, it directly affects host NIC reassembly buffer size and complexity. Finally, an $O(1)$ queue is valuable in itself, since it makes performance predictable.

Square-root growth. We find that in algorithms based on random spraying, both uplink and downlink queues grow as the square root of the message size. The proof, presented in Appendix B, is based on the fact that the number of packets that land on a given uplink or downlink is a sum of independent random variables, each variable corresponding to whether a packet lands on that link. It finds that the sum has a standard deviation of $\sqrt{m}$, which translates to a queue growth of $\sqrt{m}$. For downlinks, the proof relies on an approximation that the variations in uplink queue delays are negligible. Later, we show that simulations confirm the square-root growth (Fig. 4(b)).

Theorem 1. RSQ *and* RPS*, which are based on random spraying, see both their uplink and downlink queues grow as the square root of the message size:* $q_{\uparrow}(m) = \Theta(\sqrt{m})$ *and* $q_{\downarrow}(m) = \Theta(\sqrt{m})$ *(neglecting upstream queueing delays). Hence*

$$q(m) = \Theta\left(\sqrt{m}\right) \tag{1}$$

As we saw in Fig. 3, some destination-oblivious schemes, such as DRILL and AR, perform well locally on the uplinks, while others perform more poorly. However, all destination-oblivious schemes suffer from the downlink problem. Intuitively, it is because they simply do not optimize for this downlink and pick uplink ports independently of the exact destination of the packets. We later show that they all exhibit queues that grow with the message size, often in what appears as the square root of m (Fig. 4(b)).

Linear growth. While we proved above a square-root queue growth with the message size, we can also prove a worse queue growth for two more destination-oblivious algorithms (proof in Appendix C):

Theorem 2. SIMPLE RR *and* JSQ *keep their uplink queues bounded,* $q_{\uparrow}(m) = \Theta(1)$*, but suffer from a linear growth of their downlink queues,* $q_{\downarrow}(m) = \Theta(m)$*. Hence*

$$q(m) = \Theta(m) \tag{2}$$

Thm. 2 is the more surprising result, especially for JSQ, which is often presented as near-optimal [27, 34, 63, 64]. The cause is *synchronization*, and it is a direct consequence of the workload uniformity described in §2: all senders transmit the same message and packet size, paced at identical rates. This uniformity creates per-sender "stickiness," in which successive packets from the same sender are forwarded out the same port to the same aggregation switch. For example, consider again an edge switch whose $w = k/2$ flows all go northbound through its w ports, and assume SIMPLE RR. Since all flows are paced at the same rate, their packets arrive in the same order every period. So if the pointer sends a packet of some flow to port i, the pointer advances exactly w times before that flow's next packet, and returns to port i: the flow ends up sticking to port i for the whole run.

Given the fat tree, such sticky packets necessarily arrive at the same aggregation switch in the destination pod, then stick to the same path all the way down—where they collide with other sticky flows. Note that local *salting* does not help: even if each switch starts SIMPLE RR at a random port and proceeds in a different periodic order, the same synchronization arises. We discuss this in detail in Appendix C.

3.4 Destination-Based Rotation

Given these results, we looked to the broader scheduling literature for approaches that reach $O(1)$. We find inspiration in the DRB scheduler of Cao *et al.* [16], whose key idea is that each source host load-balances packets in round-robin fashion on a *per-destination* basis. We generalize this to a scheduling discipline we call *Destination-Based Rotation* (DR). We later contribute a switch-based DR scheme in §4.

DRB as host DR. DRB maintains round-robin pointers at each source host, one per (destination host, quantized packet size) pair, using two quanta to separate data and ACK packets. For each (host, size) pair it finds the lowest common switch layer in the topology and round-robins among the switches at that layer with a single pointer. For example, each colored source host in Fig. 2 sends its data packets to the four cores in round-robin order; this is done using encapsulation. Under failures, DRB assumes hosts have full topology knowledge and round-robins only over switches reachable by both endpoints.

Bounded queues. Between source and destination, DRB flows are load-balanced equally and periodically across all allowed paths, so any given path sees a sum of many periodic flows with random phases. We can therefore model DRB's queues with the $ND/D/1$ model [58,76], which is particularly noteworthy for *its bounded average queue size even when the arrival rate equals the service rate* (this does not happen in M/M/1 for example). We prove the following result in Appendix D. Note that the proof here is weaker than for Thms. 1 and 2, as we prove it for the DRB $ND/D/1$ approximation, not for the exact algorithm.

Theorem 3. *The average queue size for the* DRB *model is bounded:*

$$q(m) = \Theta(1). \tag{3}$$

ETHEREAL. A second destination-aware design, ETHEREAL [1], splits each flow into many subflows. Specifically, under a permutation, each host has a single flow, which ETHEREAL splits into as many equal-size subflows as there are

paths, e.g., one subflow per core switch for an inter-pod destination. Since the NIC serves its queue pairs in round-robin, the host effectively rotates its packets across all these paths, much like DRB. Furthermore, ACKs follow the reverse path of their data packets, so they are also equally load-balanced across all paths.

Thus, while ETHEREAL is not a DR scheme, since it keeps no rotation pointers, its per-destination subflows achieve a similar effect, and we obtain the same bounded-queue result (proof in Appendix E):

Theorem 4. *The average queue size for the* ETHEREAL *model is bounded:*

$$q(m) = \Theta(1). \tag{4}$$

Deployment challenges. Both schemes ask for something hosts do not have today, exhibiting significant deployment challenges. DRB requires each host to know the complete set of currently reachable switches, which in practice means running the fabric's routing protocol (e.g., BGP) on the NIC. Exposing routing and link-state updates to hosts violates the layering that separates them from the fabric's internals, and requires every NIC to maintain global topology state and re-converge on every failure. ETHEREAL requires the NIC to maintain a very large number of queue pairs (QPs) per flow (one per core, i.e., on the order of 1K with 64-port switches) in order to place each subflow on its own path. Increasing the number of QPs by 3 orders of magnitude could prove daunting at scale. Both limitations are consequences of running per-destination load balancing *at the host*. A host sees only its own flows and has no native view of the fabric. As we show next, a switch sees both.

3.5 Validation via Simulation

Since our theoretical analysis makes a number of simplifying assumptions, we present simulation results showing both the CCT and the max queue size for the various algorithms discussed above. For the CCT metric, to focus on the variable part impacted by load-balancing, we compute the CCT increase beyond a lower-bound baseline that we derive for the CCT of a permutation (Appendix A), and use naive lower bounds for non-permutation workloads. For completeness, we also show the performance of POFAN and OFAN, which are discussed later.

Fig. 4 shows that max queue size correlates well with the CCT metric from our earlier sections. We also see the general trends revealed by our theoretical analysis. Specifically, algorithms cluster in three groups: (1) SIMPLE RR and JSQ (together with CONGA) form a first group with near-linear growth of the max queue size; (2) The second group includes RPS, RSQ, REPS, and AR whose queue growth is closer to the square-root of the message size, and (3) ETHEREAL, DRB, 4OFAN and OFAN appear to reach an optimal $O(1)$ queue size.

Given these results, we next take a closer look at practically realizing DR-based LB. Since DRB and ETHEREAL face significant deployment challenges, and OFAN outperforms them in our simulations, we focus on switch-based AR and host-based REPS as leading alternatives in the remainder of the evaluation (except for Fig. 6).

4 OFAN

4.1 Algorithm

§3.4 left us with a discipline that is optimal but has only ever been realized at the host, where it is impractical. This section makes the case that the switch, rather than the host, is the natural home for DR, and presents OFAN, a switch-based realization of it.

Switch benefits. We find that the switch enjoys two significant benefits when implementing DR. First, the switch already knows the topology. It knows through BGP which links are currently alive. DRB needs the host to obtain this, e.g., by implementing somehow BGP onto the NIC; ETHEREAL keeps iterating to guess it from the host's viewpoint. A switch gets it for free, and re-converges quickly on failure through the mechanisms it already runs.

Second, and less obviously, the switch can see more traffic. This enables it to apply round-robin algorithms more efficiently, as described below.

Overview of OFAN. OFAN[2] is a switch-based realization of DR scheduling. Its goal is to maintain per-destination pointers at edge and aggregation switches, and use these pointers to select the outgoing ports for northbound packets. However, there are many possible destination hosts, leading to an impractical number of pointers to implement.

To address this scalability challenge, OFAN exploits the structured nature of a fat-tree topology to *consolidate* destinations based on their location in the tree, such that a switch only maintains a single pointer for groups of destinations that share a common edge or aggregation switch. The idea is that a switch does not need to care about the exact destination of its northbound packets, but can consolidate these packets based on their common last-mile switches.

Specifically, an edge switch E will maintain one pointer for each group of destinations that share an edge switch F ($F \neq E$). Similarly, an aggregation switch A maintains one pointer for each group of destinations that share an aggregation switch B ($B \neq A$). Hence, an edge switch on the northbound path will round-robin all traffic bound for the same destination edge switch, and an aggregation switch will round-robin traffic bound for the same destination aggregation switch.

OFAN's consolidation of destinations significantly reduces the number of pointers, while retaining DR's property of bal-

[2] OFAN's name is based on Ezekiel's vision of wheels within wheels, fitting the many pointer rotations.

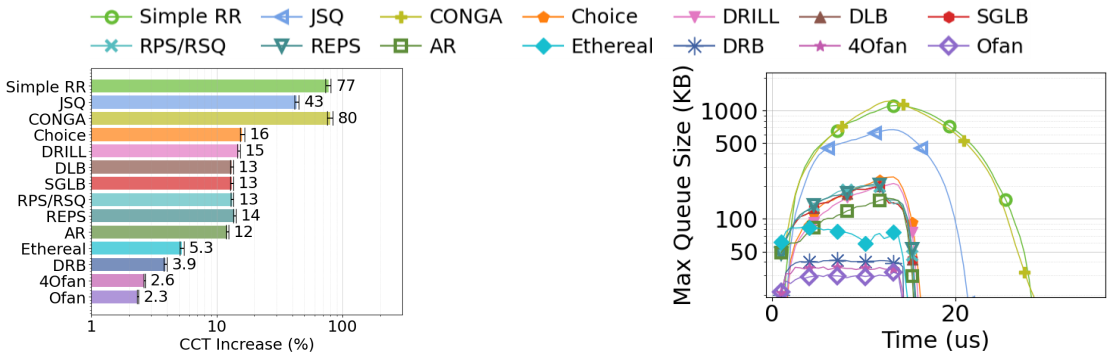

(a) CCT increase **(b) Max queue size**

Figure 4: Comparison of packet-based LB algorithms.

ancing traffic on both uplinks and downlinks. To see this, consider two flows that originate at source hosts under source edge switch E, and are destined to destination hosts under destination edge switch F. Then E and F are *mandatory waypoints* in their path. Therefore, E can maintain a single pointer that consolidates both flows. Consolidation is even more effective at the aggregation layer due to the properties of three-level fat-trees. In a fat-tree, all packets destined for a pod and come through a source aggregation switch in their uplink path can only go through a single aggregation switch in their downlink path.[3] This destination aggregation switch is again a mandatory waypoint, hence a source aggregation switch can consolidate its scheduling of all packets destined to this pod.

Per-destination pointers. The core idea in OFAN is to have *a pointer per destination*, and each incoming packet that goes up in the fat tree accesses its specific pointer to determine its northbound port. For the consolidation explained above, we want to have at each source edge switch a different pointer per destination edge switch, and at each source aggregation switch a different pointer per destination pod. Assume that each such switch has $w = k/2$ northbound ports. At start up, each pointer is initialized to a random value in $\{0,\ldots,w-1\}$. Then, at each packet arrival for this pointer, it is incremented (modulo the number w of uplink ports).

Random port order. A simple way is for each packet to be sent to the port indicated by the pointer. For example, if the pointer for the packet's destination indicates 3, then it will be sent to port 3. However, to avoid a possible synchronization across pointers, we don't want all of them to go RR in the same order. Instead, we create a random port traversal order by mapping each pointer value to an uplink port through a pointer-dependent permutation. That is, even though each pointer goes cyclically through the values $\{0,1,2,\ldots,w-1\}$ in the same RR order, its permutation will then respectively map them to some port order like $\{1,w-1,2,\ldots,0\}$. The easiest way to implement these random per-pointer permutations would be to store all of them. However, this would need a large storage space. Instead, we use a simple formula that takes the pointer value and pointer ID and returns the corresponding port ID using a XOR and an addition modulo w (see §F.2 for details).[4]

ACKs. In addition, using the same pointer for data and ACK packets would not load-balance the load equally, which is our goal: e.g., we could send a data packet on one port, then an ACK packet on the next port, etc. Thus, we follow DRB and distinguish pointers for different packet sizes. Specifically, in OFAN, an edge switch maintains a pointer for each (destination edge switch, quantized packet size) pair and an aggregation switch maintains the same for each (destination pod, quantized packet size). In our simulations, we use two quantization groups for packet sizes: one for data packets, with fixed frame sizes for data payloads, and another for ACKs. Other datacenters and workloads might require more fine-grained quantization. We initialize data pointers at random, and ACK pointers half a rotation away (offset $\lfloor w/2 \rfloor$).

Memory needs. Thus, an edge switch maintains O(# edge switches) pointers, and an aggregation switch maintains O(# pods) pointers. For example, in a fat-tree of parameter $k = 64$ with $65K$ GPUs [12, 30, 51, 66], each edge switch would maintain $4K$ pointers of $\log_2 w = 5$ bits each, totalling 2.5 KB, and each aggregation switch 126 pointers of 5 bits each, i.e., 79 B. We add 20 B to store the same basis permutation in each switch. Our discussions with switch vendors suggest that this is a very reasonable number of pointers and storage amount for switches to maintain.

Note that unlike DRB, OFAN does not require any tunneling or encapsulation of traffic, nor any change at endhosts. Of course, it does require changes at edge and aggregation switches.

Failures. When there are link failures in the topology, the

[3] To see this, notice that any packets leaving the first black aggregation switch in Fig. 2 can only traverse the second black aggregation switch on their southbound path to pod 2.

[4] As an order of magnitude, the randomized port order reduces CCT by $\approx 0.2\%$ for permutation workload, but increases it by $\approx 0.1\%$ for all-to-all; so its main appeal is in avoiding or reducing unforeseen synchronization, not necessarily in the immediate CCT reduction.

pointers still cycle through the ports in the same way. However, if they pick an invalid port for a packet, i.e., a port through which there is no possible path to the destination, then they just keep incrementing until yielding a valid port. [5]

4.2 Bounded queue sizes

Properties. We get the same $\Theta(1)$ result for OFAN's model as for DRB (again, we only prove it for the queueing model, see Appendix F):

Theorem 5. *The modeled average queue size for* OFAN *is bounded:*

$$q(m) = \Theta(1). \tag{5}$$

Consolidation pays off. Fig. 3 in §3.2 already showed OFAN flat across all four link stages. It also showed the effect consolidation buys over DRB: DRB is perfectly balanced for large messages but degrades for small ones, because its many per-flow pointers spread traffic too thin, leaving a constant residual imbalance that a small message cannot amortize. OFAN's far smaller number of consolidated pointers avoids this. In §6 we show that the gap widens with network size, tracking the widening disparity in pointer counts.

Impact on CCAs. Bounded queues change how we think about CCAs. For example, they allow the design of more aggressive CCAs with *shorter queue targets* that impact less delay-sensitive traffic (mice) (see MSwift vs. UEC's NSCC in §6).

Impact on packet sizes. Second, and surprisingly, they change the *optimal packet sizes* in the datacenter networks. The impact of packet size as a parameter in AI training is one that, to our knowledge, has not been extensively explored in the networking literature. Intuitively, sending D bytes with a packet size P and header H, involves a trade-off: large packets increase queue size and hence CCT, while small packets incur higher header overheads. We find that we can model the optimal packet size for some of our LB approaches. For DR schemes such as OFAN, if we model the $O(1)$ network queueing as a constant of α packets (independent of the number of sent packets), then we obtain the following (Appendix G):

Theorem 6. *To minimize the CCT, the optimal packet size should carry a payload that scales with $\sqrt{D}$:*

$$P - H = \sqrt{\frac{H}{\alpha} \cdot D} \tag{6}$$

Interestingly, the constant H/α reflects the battle between the two types of CCT overhead: *headers* and *queueing*. We illustrate this result in §6.

[5] We assume that the operator has a procedure in place in case the switch is a dead end.

5 POFAN Pipe-Based Implementation

5.1 POFAN

Pipes. Modern switch ASICs are organized into several parallel packet-processing *pipes* (Broadcom) or *slices* (Cisco), each handling a fixed subset of the chip's ports with dedicated stateful hardware. Because ports within the same pipe pass through identical hardware execution stages, the pipeline natively guarantees atomic Read-Modify-Write access to local registers at line rate without cross-pipe locks. Hardware pipelines historically scaled around 4-port blocks per processing engine: Broadcom's Tomahawk 3 and 4 utilize 4 ports of 400G per pipeline (1.6 Tbps per pipe) [13, 83], while 51.2 Tbps generation ASICs such as Cisco's Silicon One G200 and P200 explicitly implement 16 independent slices of 4×800G ports (3.2 Tbps per slice) [79]. Thus, we focus here on pipes of four ports.

POFAN. This pipe-based organization motivates an easier-to-implement variant, POFAN: instead of a single RR pointer shared among all bottom ports for a given destination key, as in OFAN's pipeline above, each pipe keeps its own independent RR pointer per key (e.g., with 4OFAN, the RR pointer in each pipe is shared by four inputs). This sacrifices some consolidation, and hence some of OFAN's load-balancing quality, since we are effectively splitting each pointer into several independent, uncorrelated ones, similar in spirit to DRB's many per-flow pointers (§4), in exchange for a design that fits directly within existing switch architectures, with no cross-pipe combining hardware required. By default, we assume pipes of 4 ports and use 4OFAN, matching current switch ASICs; we explore the impact of this choice on CCT performance in §6.

Example. For our running example ($k = 64$, so $w = k/2 = 32$ downlink ports) with 8 pipes of 4 ports each, POFAN replaces each of OFAN's $N \approx 4K$ shared pointers with 8 independent, per-pipe pointers, for a total of $8 \times 4K \approx 32K$ pointers in the switch.

Bounded queue sizes. POFAN keeps the same bounded queueing model as DRB and OFAN (proof in Appendix H):

Theorem 7. *The modeled average queue size for* POFAN *is bounded:*

$$q(m) = \Theta(1). \tag{7}$$

The intuition is that splitting a pointer into several independent ones keeps each of them balanced, so the load per destination on every link stays balanced as well. The price is a larger number of pointers, and hence a larger queue bound constant.

5.2 P4 implementation

We implement POFAN using Intel P4 Studio [20] on the Intel Tofino 1 platform [19] to showcase its practicality. Each pipe

Table 2: Resource consumption of POFAN on an Intel Tofino 1 switch.

Resource	Usage	
	Absolute	Relative
Match Crossbar	16 bytes	3.7%
Meter ALU	1 ALU	16.7%
Gateway	2 gateways	3.6%
SRAM	207,530 bytes	6.7%
TCAM	3,754 bytes	5.6%
VLIW Instruction	3 slots	5.6%
Hash Bits	82 bits	11.0%

stores (#edge switches · #size classes) RR pointers in a single register. The register is indexed by the destination edge switch and size class. Its value at the index is read and updated for each packet.

Table 2 summarizes the resource consumption of POFAN on an Intel Tofino 1 switch assuming a network with 64 edge switches and 4 pipes. Most of the SRAM is not used by the RR pointers themselves, but by the table that maps each (destination address, size class) pair to its RR pointer index. This table could be removed by assuming that the IP address space follows the physical topology, e.g., with a prefix per edge switch, which our implementation does not assume. The RR pointers themselves grow with the number of pipes and of edge switches, but only in steps, since Tofino allocates SRAM in blocks: the block holding the pointers for 64 edge switches holds them for up to about 1,024 edge switches, and the 2,048 edge switches of our $k = 64$ running example need a single additional block. Since combining several memory operations in a single register is common in Tofino programs, the growth remains at most linear. POFAN thus still leaves plenty of resources for the implementation of other data plane functionality.

Due to resource limitations, we cannot evaluate a hardware cluster of Tofino switches implementing POFAN. We therefore resort to simulation for performance comparison with Ofan in §6.

6 Evaluation

We now determine the robustness of our results when varying several scaling parameters (§6.1), introducing link failures (§6.2), changing the topology (§6.3), deploying 4OFAN incrementally (§6.4), using congestion control (§6.5) and considering MoE traffic (§6.6). Finally, we present results measuring CCT in a scenario simulating the training of Llama-3 [55] models with up to 405B parameters parallelized using FSDP [75, 87] and running a full stack of CCA, loss recovery, and load balancing (§6.7).

Simulator. We compare different load balancing schemes with simulations on *htsim* [11], which we chose due to its ability to scale to large datacenter networks and its existing implementation of many design points in this space. Specifically, unless stated otherwise, we simulate a 432-node 3-tier fat-tree network corresponding to $k = 12$-port switches, with 800 Gbps links, 0.5 μs link latency, RoCE-v2-sized packets with 4,096 B payloads and 62 B headers, 64 B ACKs, and 20 B gaps between frames (12 B inter-frame gap and 8 B preamble and start frame delimiter), 800 KB per-port buffers,[6], and 80 KB ECN thresholds (corresponding to 10% of the buffer size; we increase this threshold to 10 packets when studying 64 KB packets). We also perform a sensitivity analysis for many of these parameters. We run all considered scenarios at least 10 times, and display error bars for the SEM (unless they affect readability).

Workloads. In our evaluation of LB schemes, we consider an all-to-all (ATA) traffic matrix and a random permutation matrix. We focus on these two traffic matrices since they are representative of many collective algorithms, rather than modeling specific implementations of every collective. Permutation matrices are often used to implement collectives such as AllGather and AllReduce [15, 31, 80, 84] and can be used to iteratively perform an AlltoAll ($n-1$ permutation matrices iteratively). Similarly, an all-to-all traffic matrix may implement AllGathers or a one-shot AllReduce, as well as an AlltoAll. Notably, in most training jobs, multiple forms of parallelism will be used and thus no one collective will span the whole network. We explore this more in §6 and for now make the simplifying assumption of one collective at a time across all nodes. We use a default permutation message size of 1 MB.

We ran tests with the entire set of our LB schemes from §3.1. Across our tests, we find that the *relative* ranking and broad assessment of different schemes that we have made so far remains unchanged. For legibility, in this section, we compare results for switch-based AR and host-based REPS, representing state-of-the-art algorithms, against our OFAN and 4OFAN algorithms.

Metric. Our main metric is the *collective completion time* (CCT), which is the time from when the first packet of a collective is sent until the last packet of that collective is acknowledged.

6.1 Scaling

Network Size. Fig. 5 shows the impact of scaling the number of nodes in the network (and thus participating in the collective). For ATA, we illustrate two ways of scaling the network: either with constant per-flow size (Fig. 5(b)), corresponding

[6]This value roughly follows the per-port per-Gbps buffer size of modern switches [9, 67], is above a BDP, and resulted in no loss in most evaluated cases. We only use it in the CCA cases, else we use a practically infinite buffer.

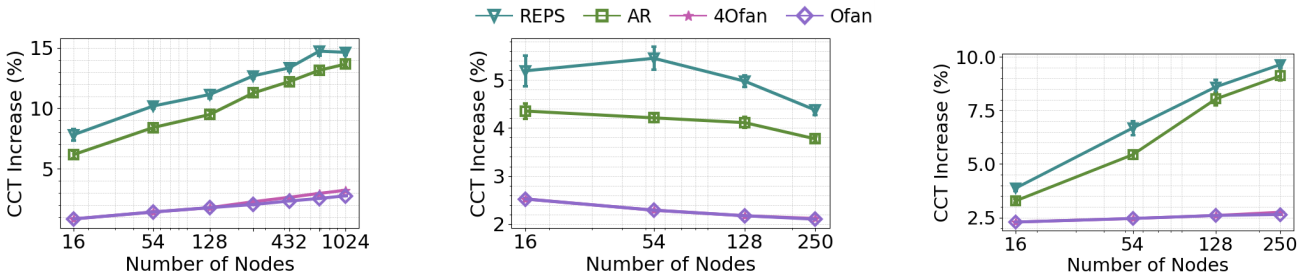

(a) Permutation **(b) ATA, constant per-flow size** **(c) ATA, constant per-host size**

Figure 5: CCT increase as network size increases.

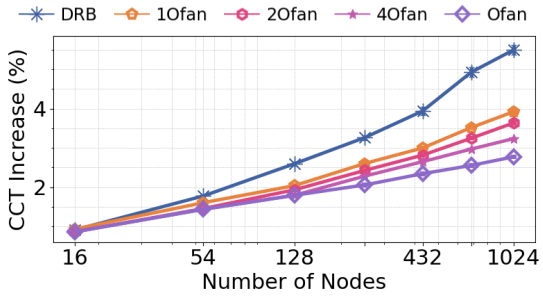

Figure 6: CCT increase of POFAN variants vs. DRB as network size increases for a random permutation.

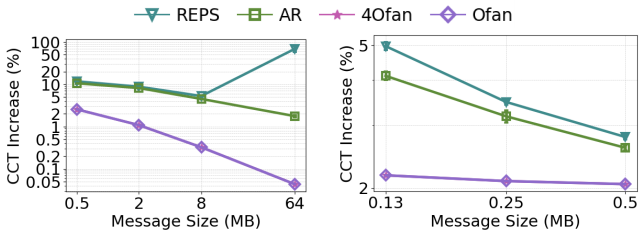

(a) Permutation **(b) All-to-All**

Figure 7: CCT increase as message size increases.

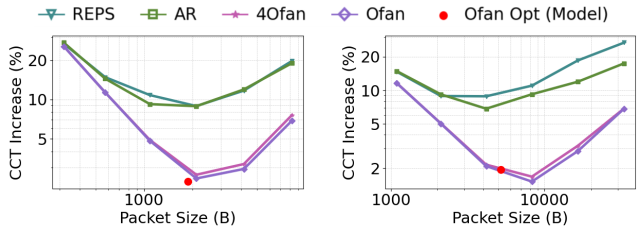

(a) 0.5 MB **message** **(b)** 4 MB **message**

Figure 8: CCT increase as packet size increases, for two message sizes in bytes.

to a model that scales with the network; or constant total host sending size (Fig. 5(c)), corresponding to adding more nodes with a constant model. We see that both the relative ordering of the schemes and the magnitude of their CCT performance differences holds or even increases as the network size grows, both for permutation and all-to-all traffic. In particular, in the permutation case, the CCT gap between 4OFAN and existing algorithms grows from about 5% with 16 nodes to about 11% with 1,024 nodes; and the gap between 4OFAN and OFAN remains reasonably small.

Pipe Size. Fig. 6 shows the impact of scaling both the pipe size in POFAN and the network size. As network grows, the gap between POFAN and OFAN predictably expands, but somewhat moderately. In addition, it is interesting to compare 1OFAN to DRB: 1OFAN holds a per-destination decision at each switch input, while DRB holds a per-destination decision at each sending host. Since each edge switch input corresponds to one host, we would expect them to be similar. However, the plot for 1OFAN exhibits and quantifies the power of consolidation at the aggregation switches, which seems to roughly correspond to half the CCT gap between OFAN and DRB.

Message Size. Fig. 7 plots the CCT increase as the message size increases. First, the relative ordering of the algorithms does not change, and 4OFAN and OFAN keep beating REPS and AR. Second, the absolute CCT gap between them decreases with increasing message size. This is because, as the message size m increases, the CCT gap is expected to increase as $\sqrt{m}$ (§3.3) while CCT grows linearly with m, so the ratio should decrease like $\sqrt{m}/m = 1/\sqrt{m}$. Third, Fig. 7(a) illustrates the surprising (and apparently previously-unreported) phenomenon that REPS gets notably worse at very large message size with permutation traffic. We conjecture that this is due to the "freezing mode" in REPS triggering as queues become long enough to cause a timeout [8].

Packet Size. We illustrate with a red dot in Fig. 8 the optimal size and modeled CCT increase vs. our best constructed lower bound. The graph shows a clear "sweet spot" in packet size at which CCT is optimized. We can see how the model roughly predicts the packet-size sweet spot observed in simulation and its CCT increase. Fig. 8(b) shows how the shorter queue sizes of OFAN also encourage larger packets: We can derive that when using spraying algorithms with square-root queue-size growth, optimal packet size only grows as $\Theta\left(D^{1/3}\right)$ rather than $\Theta\left(\sqrt{D}\right)$. Also note that we neglect second-order factors, e.g., the CRC may need to be a bit longer to protect longer packets. This result suggests many avenues for future work on

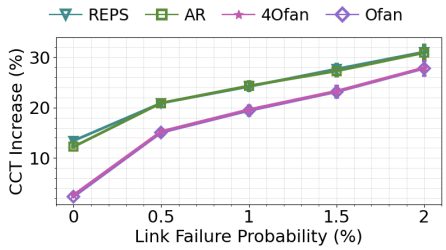

Figure 9: CCT increase due to random link failures (beyond the CCT baseline in a fully-connected network).

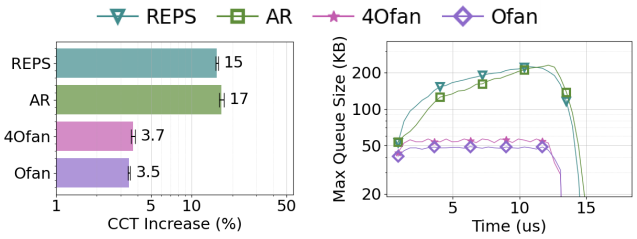

(a) CCT **(b) Max queue size**

Figure 10: Two-level fat-tree with 2,048 hosts.

optimizing performance through dynamically tuning packet size, and exploring standardization support for larger Ethernet MTUs.

6.2 Link Failures

To evaluate the impact of random link failures, we drop each edge-aggregation and aggregation-core link with a given probability in each topology instance. We then look at the increase in the average CCT, compared to the CCT lower-bound baseline without link failures. Note that in each topology instance, given the set of link drops, each host paces at a fixed rate, ρ_{max}, computed as the highest rate at which all flows could send if traffic were balanced perfectly across all valid shortest paths, to make sure it's feasible (§2.2).

Fig. 9 shows the results. We can see that the relative ordering of the LB schemes is kept about the same, and OFAN keeps outperforming REPS and AR.

6.3 Different Topologies

Fig. 10 illustrates the results of running the algorithms in a two-level fat-tree with switch radix $k = 64$ and $k^2/2 = 2,048$ hosts. OFAN and 4OFAN keep outperforming REPS and AR and reduce the CCT by over 11%.

6.4 Incremental Deployment

We want to study the impact of incrementally deploying 4OFAN in a datacenter using AR, since operators roll out new switch generations incrementally [69, 86] and may worry

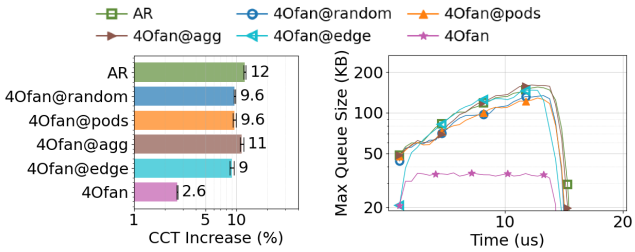

(a) CCT **(b) Max queue size**

Figure 11: Incremental deployment: scenarios where half the switches use AR and half 4OFAN.

about mixing switches with conflicting LB policies. We assume that exactly half of the LB switches (i.e., edge and aggregation) use 4OFAN, while the other half keeps using AR. We compare four ways of selecting the switches to deploy 4OFAN: (i) in random switches, (ii) in random pods (which seems closest to how operators deploy new switches, one aggregation block at a time [69]), and in (iii) all aggregation or (iv) all edge switches.

Fig. 11 conveys three lessons. First, introducing 4OFAN in some switches only decreases the CCT, i.e., mixing the two LB policies in the network is not harmful. Second, the reduction in CCT is less than proportional to the fraction of upgraded switches: the midpoint between the CCT increases of AR and 4OFAN is 7.3%, and the CCT increase in all mixes is slightly higher. This was expected, since the CCT is set by the slowest flows, which are likely to cross at least one AR switch, so measuring the worst-case scenario is intrinsically unfavorable to mixes. Last, interestingly, the best place to deploy 4OFAN is at the edge switches.

6.5 Congestion Control

We now compare per-packet spraying algorithms with CCA. We rely on three different CCAs: (i) DCTCP, a standard datacenter CCA, and two recent CCAs designed for spraying: (ii) UEC's NSCC [18, 81] and (iii) MSwift [32]. We keep evaluating a permutation traffic, but increase the message size from 1 MB to 4 MB to give enough time for the CCAs to adapt (propagation RTT is about 150 packets, so we increase the message size from 1.7 RTTs to 6.8 RTTs).

Fig. 12 illustrates the CCA impact. With DCTCP, the CCTs of REPS and AR are just awful, as they are multiplied by 3–4×, while POFAN and OFAN have a CCT increase under 1%. The main reason is that even though htsim's default DCTCP implementation allows for high duplicate-ACK threshold (specifically, 3 plus the number of paths that can be taken by the flow), a few flows still reach the threshold, aggressively decreasing their windows and hurting the CCT. With such a stark difference, we drop DCTCP from now on.

With the spraying-based CCAs, REPS and AR also perform worse, but by more reasonable margins. Each CCA needs to rein down the flow rates for REPS and AR, and

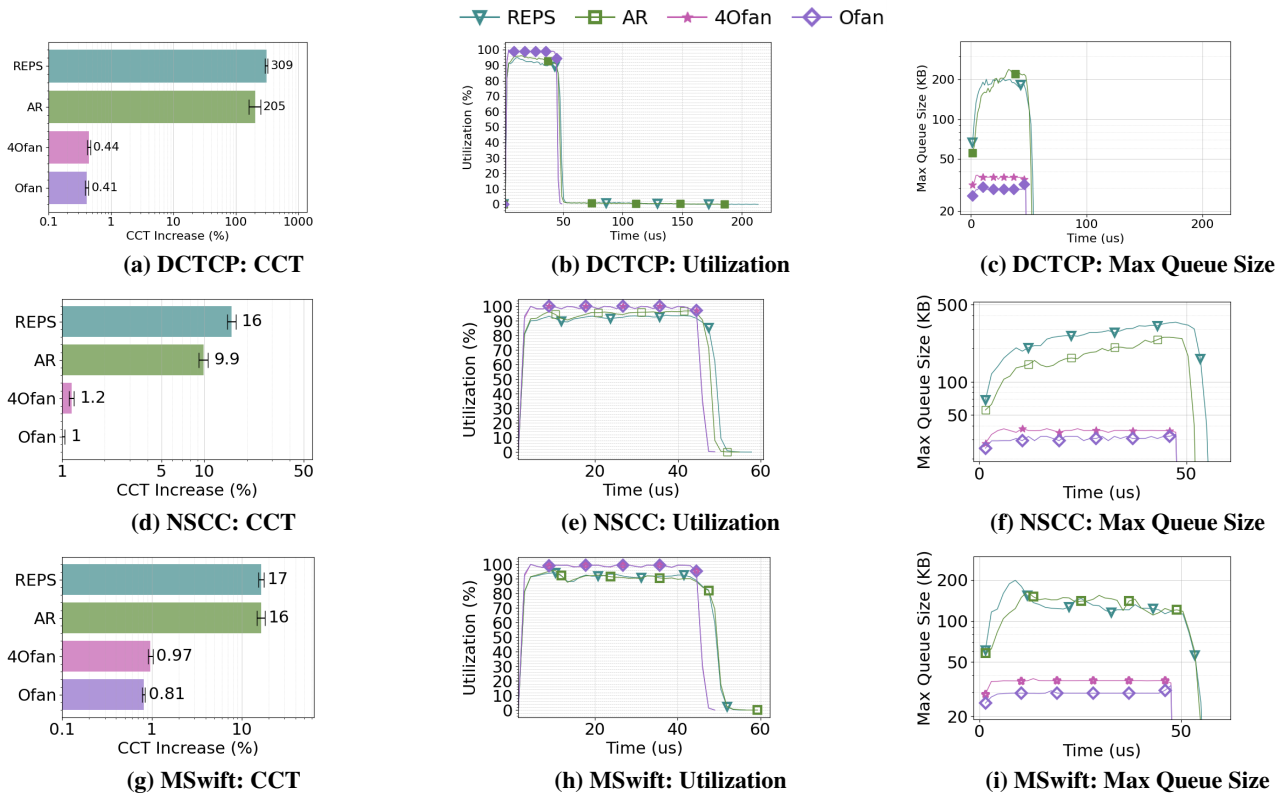

(a) DCTCP: CCT (b) DCTCP: Utilization (c) DCTCP: Max Queue Size

(d) NSCC: CCT (e) NSCC: Utilization (f) NSCC: Max Queue Size

(g) MSwift: CCT (h) MSwift: Utilization (i) MSwift: Max Queue Size

Figure 12: Permutation with 4 MB messages using either (a-c) DCTCP, (d-f) UEC's NSCC or (g-i) MSwift.

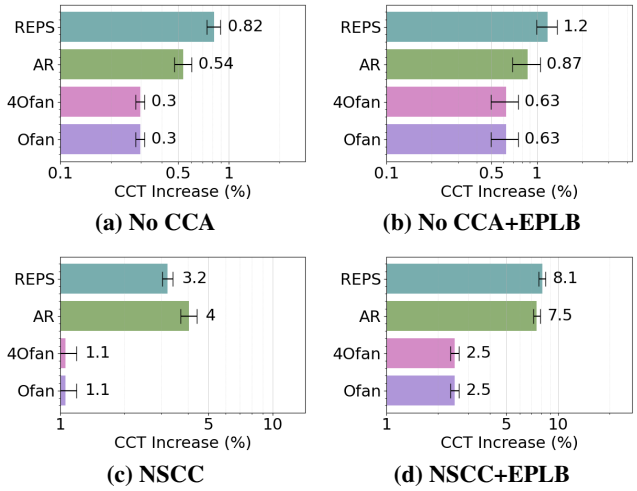

(a) No CCA (b) No CCA+EPLB

(c) NSCC (d) NSCC+EPLB

Figure 13: MoE workload: (a-b) without CCA and (c-d) with NSCC, either without or with EPLB reducing the skew.

therefore their CCT inflation can worsen by at least 8% with NSCC and 14% with MSwift vs. OFAN and 4OFAN, which do not need to be slowed down as much. The difference is a bit larger with MSwift as it targets smaller queues than NSCC.

6.6 MoE Workload

We generate an MoE routing trace by running one training iteration of the open-source Qwen3-235B model, which has 128 experts per layer and activates 8 experts per token [73], starting from its released weights with synthetic input tokens, and recording the token-to-expert routing decisions of one layer. The trace yields a 128×128 GPU-to-expert matrix of token counts. Each token copy carries 8,192 B for a hidden size of 4,096 in BF16. We then consider a $k = 8$ fat tree with $n = 128$ GPUs. At each run, we randomly map one expert per GPU, and obtain a 128×128 GPU-to-GPU traffic matrix. As in DeepSeek-V3's DeepEP [23, 54], dispatch is a hierarchical all-to-all over 8-GPU servers: each GPU sends the tokens for a remote server to its same-rank peer there, which forwards them over NVLink, so only inter-server traffic crosses the fat tree. Each GPU sends 492 MB on average over the fat tree, and the load is skewed: the busiest GPU receives $1.48\times$ the mean (728 MB). Given these large messages, we use 64 KB packets, in line with the larger optimal packet size of Thm. 6, which also keeps the simulations tractable.

To evaluate the sensitivity to a less skewed MoE workload, we also replicate popular experts using expert-parallel load balancing (EPLB) [24, 54], such that each GPU hosts two expert slots, which reduces the skew to $1.13\times$.

Fig. 13 illustrates the results. We start with an idealized example without CCA (Figs. 13(a) and 13(b)), where we just

allow all sources to blast traffic with infinite buffers and no losses. Since MoE traffic is incast-like and is bound by the most requested GPU, the bottleneck is at the receiver and not in the network. Thus, buffers grow rapidly (beyond 100 MB), and the LB impact is small, although OFAN's CCT is still a bit better.

We next employ NSCC, and differences become larger, especially with a smaller skew (Fig. 13(d)), as LB takes a larger role. OFAN clearly outperforms other algorithms. Note that since $w = k/2 = 4$ here, OFAN and 4OFAN are equivalent.

6.7 Training Scenario

Finally, we evaluate the performance of each LB scheme on a realistic AI training scenario. Specifically, we simulate training Llama-3 [55] (8B, 70B, and 405B) parallelized using Fully Sharded Data Parallel (FSDP2) [29, 75, 87] with HSDP configuration on a 1,024-GPU cluster. The cluster is organized as 128 servers of 8 GPUs, and we assume random server placement. We evaluate the main network bottleneck in FSDP2 training by focusing on the backward pass, as it combines a ReduceScatter of gradients and an AllGather of weights. FSDP2 traffic is implemented as a hierarchical ring, where intra-server communications are performed in the high-bandwidth domain (scale-up, like NVLink), while our fat-tree network implements inter-server communications (scale-out) using 8 parallel rings: each logical GPU i is connected to logical GPU $i+8 \pmod{1,024}$. Because sharding is 1,024-way while each ring spans only 128 GPUs, a ring step carries one $1/1{,}024$ shard rather than $1/128$ of a layer; the intra-server AllGather then completes the remaining $7/8$. We assume FP8 precision and 4 KB packet payload, translating to 119 packets per flow per layer in Llama-3 8B (32 layers), 418 in Llama-3 70B (80 layers) and 1533 in Llama-3 405B (126 layers). We enable both the NSCC and MSwift CCAs with loss recovery, and measure the CCT of each collective phase.

As shown in Fig. 14, while REPS and AR overheads remain about constant, the relative CCT increase for OFAN drops as model size scales. When compared against the best out of the two existing models, at 405B parameters, OFAN (resp. 4OFAN) reduces the CCT increase beyond baseline by 16× (resp. 15×) with NSCC and 39× (resp. 36×) with MSwift. This improvement stems from superior utilization (Figs. 14(d) and 14(i)): Since the CCA must throttle REPS and AR to maintain stable queue sizes (Figs. 14(e) and 14(j)), OFAN achieves significantly higher effective throughput.

7 Discussion

Realizing OFAN's pointers at line rate. §5 presented POFAN, the variant that sidesteps cross-pipe coordination by keeping an independent pointer per pipe. It is fair to ask what it would take to implement OFAN's shared pointers exactly, without that approximation. The difficulty is concurrent access: a switch processes many packets per cycle, and several of them may be bound for the same destination switch and therefore contend for the same round-robin register. This is a classical problem with many processors issuing fetch-and-add to a shared counter. Its classical solution, a *combining network* [46, 50], merges duplicate requests before they reach the counter. The general machinery is too costly and too variable in latency for a switch pipeline, but our problem is far narrower: a small, fixed set of local pointers. In Appendix I we show that a fixed-latency, four-stage pipeline suffices. It sorts the batch by pointer key, combines packets sharing a key, updates each key's register once by the group's count, then decombines and unsorts so each packet leaves with its own port. A back-of-envelope estimate puts this at roughly 10^5 gates and a few milliwatts, against a switch ASIC that draws hundreds of watts: negligible. We therefore regard OFAN's exact form as implementable, with POFAN as the pragmatic near-term option.

8 Related Work

There is a vast literature on load balancing in datacenters, and we cannot hope to summarize it in any depth. The citations in Table 1 provide a large subset. Many proposed systems [31,41,72,82] do not seriously consider packet spraying given its potential adverse effects on the transport due to reordering. Some systems do include spraying in their design [4,14,35,68], but often introduce specialized mechanisms to make spraying work (e.g., proposing an entirely new transport). Two exceptions are [26] and [8], where packet spraying is evaluated against ECMP and some other load balancing granularities, but one transport is assumed and the distance from optimality is unclear.

We already discussed throughout the paper many LB schemes that we compare against: switch- and host-based spraying in §1 and Table 1, and the destination-based DRB and ETHEREAL in §3.4.

Fat-tree routing in HPC. Routing in fat trees has long been studied for HPC interconnects [27]. Deterministic routing sends all the traffic to a given destination along a fixed path, as in ECMP, without packet spraying. In the paper we focus more on packet spraying, which can achieve better load balancing.

Load-balanced routers. Destination-based rotation has roots in three-stage load-balanced router algorithms such as UFS, FOFF and Padded Frames [39, 43, 44]. Load-balanced routers spread the traffic of each flow cyclically over intermediate ports to guarantee full throughput.

SGLB. Broadcom offers a Global Load Balancing (GLB) feature in its Tomahawk chips [59, 70]. However, GLB is limited to 3-stage Clos networks with at most 64 edge switches, and cannot run end-to-end on a 5-layer fat-tree topology [59]. This is because GLB relies on the "Next-Next-Hop" (NNH) visi-

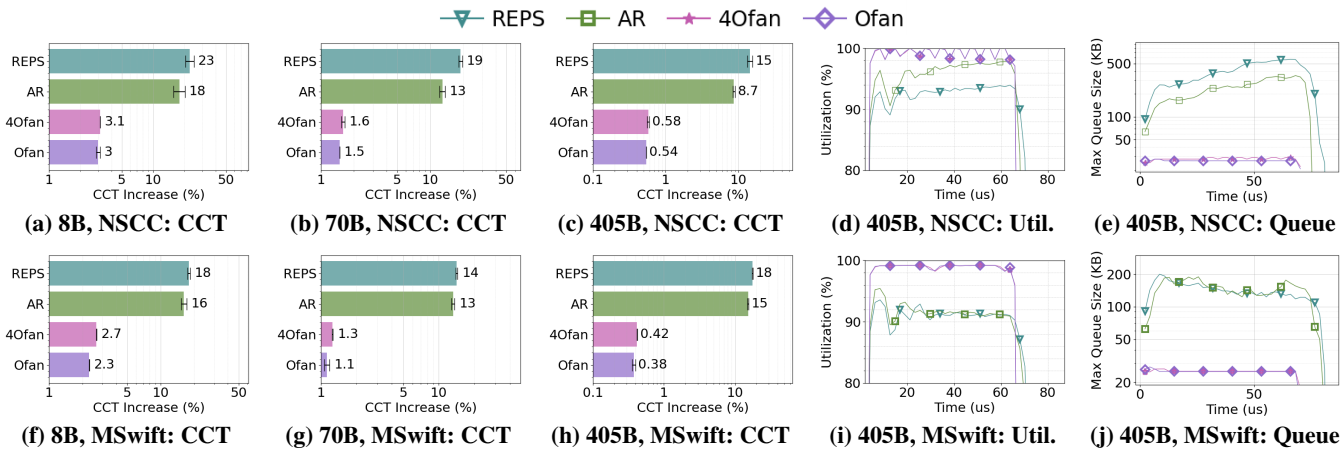

(a) 8B, NSCC: CCT (b) 70B, NSCC: CCT (c) 405B, NSCC: CCT (d) 405B, NSCC: Util. (e) 405B, NSCC: Queue

(f) 8B, MSwift: CCT (g) 70B, MSwift: CCT (h) 405B, MSwift: CCT (i) 405B, MSwift: Util. (j) 405B, MSwift: Queue

Figure 14: FSDP2 with Llama-3 8B, 70B and 405B parameters using either (a-e) NSCC or (f-j) MSwift.

bility. This inherently limits the decision scope. For example, an edge switch in a 5-layer fat-tree cannot see the congestion status of the link between a core switch and a distant aggregation switch. ByteDance's SGLB introduces several ways to mitigate this limitation in three-level fat trees [70], which we apply in our evaluations.

Acknowledgments

We thank Nandita Dukkipati, Ziming Mao and David Tennenhouse for their helpful comments. This work was partly supported by the Louis and Miriam Benjamin Chair in Computer-Communication Networks.

9 Conclusion

AI training pushes every source to 100% of line rate, leaving no slack to hide imperfect load balancing: queues grow with the message size and utilization stalls. Because a fat tree's downward path to a destination is unique, destination-oblivious schemes cannot undo the imbalance they create on the way up, as we proved for randomized spraying using an $\Omega(\sqrt{m})$ queueing lower bound.

OFAN instead rotates per destination in the switch, consolidating all traffic sharing a destination onto one pointer. Moreover, we showed how its 4OFAN variant fits commodity ASICs and cuts CCT inflation by 15–36× in an end-to-end FSDP2 evaluation at 405B parameters.

References

[1] Vamsi Addanki, Prateesh Goyal, Ilias Marinos, and Stefan Schmid. Ethereal: Divide and conquer network load balancing in large-scale distributed training. arXiv preprint arXiv:2407.00550, 2025.

[2] Ronald G Addie and Moshe Zukerman. An approximation for performance evaluation of stationary single server queues. *IEEE Transactions on Communications*, 42(12):3150–3160, 2002.

[3] Mohammad Alizadeh, Tom Edsall, Sarang Dharmapurikar, Ramanan Vaidyanathan, Kevin Chu, Andy Fingerhut, Vinh The Lam, Francis Matus, Rong Pan, Navindra Yadav, et al. CONGA: Distributed congestion-aware load balancing for datacenters. In *Proceedings of the 2014 ACM conference on SIGCOMM*, pages 503–514, 2014.

[4] Mohammad Alizadeh, Shuang Yang, Milad Sharif, Sachin Katti, Nick McKeown, Balaji Prabhakar, and Scott Shenker. pFabric: Minimal near-optimal datacenter transport. *ACM SIGCOMM Computer Communication Review*, 43(4):435–446, 2013.

[5] Joao Araujo, Alex Chow, Mark Handley, Ryder Lewis, Christoph Paasch, Jitendra Padhye, Michael Papamichael, Greg Steinbrecher, Amin Tootoonchian, Lihua Yuan, et al. Resilient AI supercomputer networking using MRC and SRv6. arXiv preprint arXiv:2605.04333, 2026.

[6] Arista. UCMP. https://www.arista.com/en/support/toi/tag/ucmp, 2026.

[7] Daniel Berend and Aryeh Kontorovich. A sharp estimate of the binomial mean absolute deviation with applications. *Statistics & Probability Letters*, 83(4):1254–1259, 2013.

[8] Tommaso Bonato, Abdul Kabbani, Ahmad Ghalayini, Michael Papamichael, Mohammad Dohadwala, Lukas Gianinazzi, Mikhail Khalilov, Elias Achermann, Daniele De Sensi, and Torsten Hoefler. REPS: Recycled entropy packet spraying for adaptive load balancing and failure mitigation. In *EuroSys*, 2026.

[9] Broadcom. Tomahawk3 / BCM56980 series. https://www.broadcom.com/products/ethernet-connectivity/switching/strataxgs/bcm56980-series.

[10] Broadcom. Trident 3 dynamic load balancing. https://www.broadcom.com/video/6328983519112.

[11] Broadcom. csg-htsim. https://github.com/Broadcom/csg-htsim, 2023.

[12] Broadcom. Tomahawk4 / BCM56990 Series. https://www.broadcom.com/products/ethernet-connectivity/switching/strataxgs/bcm56990-series, 2026.

[13] Broadcom Inc. BCM56980 data sheet: 12.8 Tb/s multi-layer switch. https://docs.broadcom.com/doc/56980-DS, 2019. Doc. 56980-DS111, July 19, 2019.

[14] Qizhe Cai, Mina Tahmasbi Arashloo, and Rachit Agarwal. dcPIM: Near-optimal proactive datacenter transport. In *Proceedings of the ACM SIGCOMM 2022 Conference*, pages 53–65, 2022.

[15] Zixian Cai, Zhengyang Liu, Saeed Maleki, Madanlal Musuvathi, Todd Mytkowicz, Jacob Nelson, and Olli Saarikivi. Synthesizing optimal collective algorithms. In *Proceedings of the 26th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, PPoPP '21, page 62–75, New York, NY, USA, 2021. Association for Computing Machinery.

[16] Jiaxin Cao, Rui Xia, Pengkun Yang, Chuanxiong Guo, Guohan Lu, Lihua Yuan, Yixin Zheng, Haitao Wu, Yongqiang Xiong, and Dave Maltz. Per-packet load-balanced, low-latency routing for Clos-based data center networks. In *ACM CoNEXT*, pages 49–60, 2013.

[17] Ren Chen and Viktor K. Prasanna. Computer generation of high throughput and memory efficient sorting designs on FPGA. *IEEE Transactions on Parallel and Distributed Systems*, 28(11):3100–3113, 2017.

[18] Ultra Ethernet Consortium. Ultra Ethernet specification v1.0.2. https://ultraethernet.org/wp-content/uploads/sites/20/2026/01/UE-Specification-1.0.2-1.pdf, Jan 2026.

[19] Intel Corporation. Intel Intelligent Fabric Processors. https://www.intel.com/content/www/us/en/products/network-io/programmable-ethernet-switch.html, accessed Apr 15, 2026.

[20] Intel Corporation. Intel P4 Studio Software Development Environment (SDE). https://github.com/p4lang/open-p4studio, accessed Apr 6, 2026.

[21] NVIDIA Corporation. Collective operations. https://docs.nvidia.com/deeplearning/nccl/user-guide/docs/usage/collectives.html, 2020.

[22] NVIDIA Corporation. NVIDIA collective communications library (NCCL). https://developer.nvidia.com/nccl, 2025.

[23] DeepSeek-AI. DeepEP: An efficient expert-parallel communication library. https://github.com/deepseek-ai/DeepEP, 2025.

[24] DeepSeek-AI. EPLB: Expert parallelism load balancer. https://github.com/deepseek-ai/EPLB, 2025.

[25] Giorgos Dimitrakopoulos, Konstantinos Papachatzopoulos, and Vassilis Paliouras. Sum propagate adders. *IEEE Transactions on Emerging Topics in Computing*, 9(3):1479–1488, 2021.

[26] Advait Dixit, Pawan Prakash, Y Charlie Hu, and Ramana Rao Kompella. On the impact of packet spraying in data center networks. In *2013 Proceedings IEEE INFOCOM*, pages 2130–2138. IEEE, 2013.

[27] Jose Duato, Sudhakar Yalamanchili, and Lionel M. Ni. *Interconnection Networks: An Engineering Approach*. Morgan Kaufmann Publishers Inc, 2002.

[28] Rick Durrett. *Probability: Theory and Examples*. Cambridge University Press, 5th edition, 2019.

[29] Wei Feng, Will Constable, and Yifan Mao. Getting started with Fully Sharded Data Parallel (FSDP2). https://pytorch.org/docs/stable/fsdp.html, 2025.

[30] FS. NVIDIA® 64-Port NDR 400G InfiniBand Data Center Switch. https://www.fs.com/products/194710.html, 2026.

[31] Adithya Gangidi, Rui Miao, Shengbao Zheng, Sai Jayesh Bondu, Guilherme Goes, Hany Morsy, Rohit Puri, Mohammad Riftadi, Ashmitha Jeevaraj Shetty, Jingyi Yang, et al. RDMA over Ethernet for distributed training at Meta scale. In *ACM SIGCOMM*, pages 57–70, 2024.

[32] Barak Gerstein, Mark Silberstein, and Isaac Keslassy. Making congestion control robust to per-packet load balancing in datacenters. *arXiv preprint arXiv:2509.07907*, 2025.

[33] Soudeh Ghorbani, Zibin Yang, P Brighten Godfrey, Yashar Ganjali, and Amin Firoozshahian. Drill: Micro load balancing for low-latency data center networks. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*, pages 225–238, 2017.

[34] Crispín Gomez, Francisco Gilabert, María Engracia Gomez, Pedro López, and José Duato. Deterministic versus adaptive routing in fat-trees. In *2007 IEEE International Parallel and Distributed Processing Symposium*, pages 1–8. IEEE, 2007.

[35] Mark Handley, Costin Raiciu, Alexandru Agache, Andrei Voinescu, Andrew W Moore, Gianni Antichi, and Marcin Wójcik. Re-architecting datacenter networks and stacks for low latency and high performance. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*, pages 29–42, 2017.

[36] Christian Hopps. Analysis of an Equal-Cost Multi-Path Algorithm. RFC 2992, November 2000.

[37] HPE Juniper. BGP Link-Bandwidth Community. https://www.juniper.net/documentation/us/en/software/junos/ai-ml-evo/bgp/topics/topic-map/bgp-link-bandwidth.html, 2026.

[38] Huawei. CloudEngine 9800, 8800, and 6800 V300R024C10, C11 Configuration Guide - IP Routing. https://info.support.huawei.com/hedex/api/pages/EDOC1100277644/AEM10221/04/resources/vrp/dc_vrp_bgp_cfg_4094.html, 2026.

[39] J.J. Jaramillo, F. Milan, and R. Srikant. Padded frames: A novel algorithm for stable scheduling in load-balanced switches. *IEEE/ACM Transactions on Networking*, 16(5):1212–1225, 2008.

[40] Chao Jin, Ziheng Jiang, Zhihao Bai, Zheng Zhong, Juncai Liu, Xiang Li, Ningxin Zheng, Xi Wang, Cong Xie, Qi Huang, et al. MegaScale-MoE: Large-scale communication-efficient training of mixture-of-experts models in production. arXiv preprint arXiv:2505.11432, 2025.

[41] Srikanth Kandula, Dina Katabi, Shantanu Sinha, and Arthur Berger. Dynamic load balancing without packet reordering. *ACM SIGCOMM Computer Communication Review*, 37(2):51–62, 2007.

[42] Naga Katta, Aditi Ghag, Mukesh Hira, Isaac Keslassy, Aran Bergman, Changhoon Kim, and Jennifer Rexford. Clove: Congestion-aware load balancing at the virtual edge. In *Proceedings of the 13th International Conference on emerging Networking EXperiments and Technologies*, pages 323–335, 2017.

[43] Isaac Keslassy. *The load-balanced router*. Stanford University, 2004.

[44] Isaac Keslassy, Shang-Tse Chuang, Kyoungsik Yu, David Miller, Mark Horowitz, Olav Solgaard, and Nick McKeown. Scaling internet routers using optics. *ACM SIGCOMM*, 33(4):189–200, 2003.

[45] Sajy Khashab, Albert Gran Alcoz, Alon Gal, Jacky Romano, Rani Abboud, Yonatan Piasetzky, Lior Maman, Amit Nishry, Barak Gafni, Omer Shabtai, et al. High-speed networking for giga-scale AI factories. arXiv preprint arXiv:2605.21187, 2026.

[46] Clyde P. Kruskal, Larry Rudolph, and Marc Snir. Efficient synchronization on multiprocessors with shared memory. *ACM Transactions on Programming Languages and Systems*, 10(4):579–601, October 1988.

[47] Jennifer Langston. Microsoft announces new supercomputer, lays out vision for future AI work. https://news.microsoft.com/source/features/ai/openai-azure-supercomputer/, 2020.

[48] Petr Lapukhov, Ariff Premji, and Jon Mitchell. Use of BGP for Routing in Large-Scale Data Centers. RFC 7938, August 2016.

[49] Yanfang Le, Rong Pan, Peter Newman, Jeremias Blendin, Abdul Kabbani, Vipin Jain, Raghava Sivaramu, and Francis Matus. STrack: A reliable multipath transport for AI/ML clusters. arXiv preprint arXiv:2407.15266, 2024.

[50] Alvin R. Lebeck and Gurindar S. Sohi. Request combining in multiprocessors with arbitrary interconnection networks. *IEEE Transactions on Parallel and Distributed Systems*, 5(11):1140–1155, 1994.

[51] Lee, Kevin and Gangidi, Adi and Oldham, Mathew. Building Meta's GenAI Infrastructure. https://engineering.fb.com/2024/03/12/data-center-engineering/building-metas-genai-infrastructure/, 2024.

[52] Shenggui Li and Siqi Mai. Paradigms of parallelism. https://colossalai.org/docs/concepts/paradigms_of_parallelism/.

[53] Stephane Litkowski, Satya R Mohanty, Arie Vayner, Akshay Gattani, Ajay Kini, Jeff Tantsura, and Reshma Das. BGP link bandwidth extended community use cases. Internet-Draft draft-ietf-bess-ebgp-dmz-08, Internet Engineering Task Force, October 2025. Work in Progress.

[54] Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. DeepSeek-V3 technical report. arXiv preprint arXiv:2412.19437, 2025.

[55] AI @ Meta Llama Team. The Llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.

[56] Jie Lu, Jiaqi Gao, Fei Feng, Zhiqiang He, Menglei Zheng, Kun Liu, Jun He, Binbin Liao, Suwei Xu, Ke Sun, et al. Alibaba Stellar: A new generation RDMA network

for cloud AI. In *Proceedings of the ACM SIGCOMM 2025 Conference*, SIGCOMM '25, page 453–466, New York, NY, USA, 2025. Association for Computing Machinery.

[57] Sarah McClure, Evyatar Cohen, Alex Shpiner, Mark Silberstein, Sylvia Ratnasamy, Scott Shenker, and Isaac Keslassy. Load balancing for AI training workloads. arXiv preprint arXiv:2507.21372, 2026.

[58] Michael Menth and Stefan Muehleck. Packet waiting time for multiplexed periodic on/off streams in the presence of overbooking. *International Journal of Communication Networks and Distributed Systems*, 4(2):207–229, 2010.

[59] Michal Styszynski. HPE Juniper: Avoiding AI/ML traffic congestion with global load balancing. https://blogs.juniper.net/en-us/industry-solutions-and-trends/avoiding-ai-ml-traffic-congestion-with-global-load-balancing, 2024.

[60] Prodosh Mohapatra, Reshma Das, SATYA R MOHANTY, Serge Krier, Rafal Jan Szarecki, and Akshay Gattani. BGP Link Bandwidth Extended Community. Internet-Draft draft-ietf-idr-link-bandwidth-24, Internet Engineering Task Force, January 2026. Work in Progress.

[61] NOKIA. BGP. https://documentation.nokia.com/srlinux/25-7/books/routing-protocols/bgp.html, 2026.

[62] NVIDIA. Out-of-order data placement. https://docs.nvidia.com/doca/archive/2-9-4/out-of-order-data-placement/index.html.

[63] NVIDIA. NVIDIA InfiniBand Adaptive Routing Technology Accelerating HPC and AI Applications. https://www.amax.com/content/files/2023/12/NVIDIA_InfiniBand_Adaptive_Routing_Technology_Insights_Whitepaper.pdf, 2023.

[64] NVIDIA. NVIDIA Spectrum-X Network Platform Architecture. https://resources.nvidia.com/en-us-networking-ai/nvidia-spectrum-x, 2024.

[65] NVIDIA. BGP Weighted Equal Cost Multipath. https://docs.nvidia.com/networking-ethernet-software/cumulus-linux-515/Layer-3/Routing/BGP-Weighted-Equal-Cost-Multipath/, 2026.

[66] NVIDIA. NVIDIA DGX SuperPOD: Next Generation Scalable Infrastructure for AI Leadership Reference Architecture Featuring NVDIA DGX H100 . https://docs.nvidia.com/dgx-superpod/reference-architecture-scalable-infrastructure-h100/latest/components.html, 2026.

[67] NVIDIA Corporation. NVIDIA Spectrum SN5000 series switches. https://nvdam.widen.net/s/mmvbnpk8qk/networking-ethernet-switches-sn5000-datasheet-us.

[68] Vladimir Olteanu, Haggai Eran, Dragos Dumitrescu, Adrian Popa, Cristi Baciu, Mark Silberstein, Georgios Nikolaidis, Mark Handley, and Costin Raiciu. An edge-queued datagram service for all datacenter traffic. In *USENIX NSDI*, pages 761–777, 2022.

[69] Leon Poutievski, Omid Mashayekhi, Joon Ong, Arjun Singh, Mukarram Tariq, Rui Wang, Jianan Zhang, Virginia Beauregard, Patrick Conner, Steve Gribble, et al. Jupiter evolving: Transforming Google's datacenter network via optical circuit switches and software-defined networking. In *ACM SIGCOMM*, 2022.

[70] Chenchen Qi, Wenfei Wu, Yongcan Wang, Keqiang He, Yu-Hsiang (Sean) Kao, Zongying He, Chen-Yu Yen, Zhuo Jiang, Feng Luo, Surendra Anubolu, et al. SGLB: Scalable and robust global load balancing in commodity AI clusters. In *ACM SIGCOMM*, page 626–644, 2025.

[71] Kun Qian, Yongqing Xi, Jiamin Cao, Jiaqi Gao, Yichi Xu, Yu Guan, Binzhang Fu, Xuemei Shi, Fangbo Zhu, Rui Miao, et al. Alibaba HPN: A data center network for large language model training. In *ACM SIGCOMM*, pages 691–706, 2024.

[72] Mubashir Adnan Qureshi, Yuchung Cheng, Qianwen Yin, Qiaobin Fu, Gautam Kumar, Masoud Moshref, Junhua Yan, Van Jacobson, David Wetherall, and Abdul Kabbani. Plb: congestion signals are simple and effective for network load balancing. In *ACM SIGCOMM*, pages 207–218, 2022.

[73] Qwen Team. Qwen3 technical report. arXiv preprint arXiv:2505.09388, 2025.

[74] Costin Raiciu, Dragos Niculescu, Marcelo Bagnulo, and Mark James Handley. Opportunistic mobility with multipath TCP. In *Proceedings of the sixth international workshop on MobiArch*, pages 7–12, 2011.

[75] Samyam Rajbhandari, Jeff Rasley, Olatunji Ruwase, and Yuxiong He. ZeRO: Memory optimizations toward training trillion parameter models. In *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–16. IEEE, 2020.

[76] James Roberts, Ugo Mocci, and Jorma Virtamo. *Broadband network traffic: Performance evaluation and design of broadband multiservice networks*. Springer, 1996.

[77] Arjun Singhvi, Nandita Dukkipati, Prashant Chandra, Hassan MG Wassel, Naveen Kr Sharma, Anthony Rebello, Henry Schuh, Praveen Kumar, Behnam Montazeri, Neelesh Bansod, et al. Falcon: A reliable, low latency hardware transport. In *ACM SIGCOMM*, pages 248–263, 2025.

[78] Cha Hwan Song, Xin Zhe Khooi, Raj Joshi, Inho Choi, Jialin Li, and Mun Choon Chan. Network load balancing with in-network reordering support for RDMA. In *ACM SIGCOMM*, page 816–831, 2023.

[79] Faraz Taifehesmatian. Nexus 9000 architecture. Cisco Live session BRKDCN-3939, 2026.

[80] Rajeev Thakur and William D. Gropp. Improving the performance of collective operations in MPICH. In Jack Dongarra, Domenico Laforenza, and Salvatore Orlando, editors, *Recent Advances in Parallel Virtual Machine and Message Passing Interface*, pages 257–267, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg.

[81] Ultra Ethernet Consortium. Ultra Ethernet. https://ultraethernet.org/, 2026.

[82] Erico Vanini, Rong Pan, Mohammad Alizadeh, Parvin Taheri, and Tom Edsall. Let it flow: Resilient asymmetric load balancing with flowlet switching. In *USENIX NSDI*, pages 407–420, 2017.

[83] Bob Wheeler. Tomahawk 4 switch first to 25.6Tbps: Broadcom doubles 400Gbps ports with unprecedented 512 Serdes. *Microprocessor Report*, December 2019.

[84] Cliff Woolley. NCCL: Accelerated multi-GPU collective communications. https://images.nvidia.com/events/sc15/pdfs/NCCL-Woolley.pdf.

[85] Hongyi Zeng, Min Si, Pavan Balaji, Yongzhou Chen, Ching-Hsiang Chu, Adithya Gangidi, Prashanth Kannan, Bingzhe Liu, Saif Hasan, Dong He, et al. Connecting 100K+ GPUs: Building the communication stack for large-scale LLM training. In *ACM SIGCOMM*, pages 505–518, 2026.

[86] Mingyang Zhang, Radhika Niranjan Mysore, Sucha Supittayapornpong, and Ramesh Govindan. Understanding lifecycle management complexity of datacenter topologies. In *USENIX NSDI*, 2019.

[87] Yanli Zhao, Andrew Gu, Rohan Varma, Liang Luo, Chien-Chin Huang, Min Xu, Less Wright, Hamid Shojanazeri, Myle Ott, Sam Shleifer, et al. PyTorch FSDP: Experiences on scaling fully sharded data parallel. *Proceedings of the VLDB Endowment*, 16(12):3848–3860, 2023.

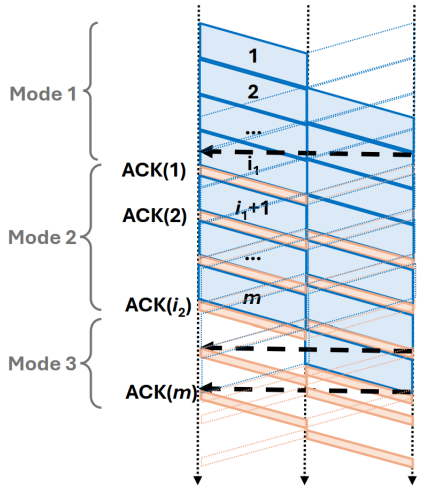

Figure 15: Illustration of the symmetric dynamics of the permutation collective.

[88] Junlan Zhou, Malveeka Tewari, Min Zhu, Abdul Kabbani, Leon Poutievski, Arjun Singh, and Amin Vahdat. WCMP: Weighted cost multipathing for improved fairness in data centers. In *Proceedings of the Ninth European Conference on Computer Systems*, pages 1–14, 2014.

A Permutation CCT Lower-Bound

Outline. Our goal is to provide a tight lower-bound baseline for the CCT of a permutation workload when hosts send at 100% of the link rate. This baseline is used in all of the figures that study the CCT increase caused by a load-balancing scheme beyond this baseline.

Example. With our default assumptions (§6), we get below a lower bound of 17.05694 μs for a message size $m = 256$. We want to compare that to smallest value that we can get. We only send a flow from the first host to the last and back in a small network ($k = 4$), using all algorithms, iterate 1000 times, and take the minimum over all possibilities, we measure in practice a value of 17.0587 μs (with $k = 8$, we get 17.0609 μs, which is less tight). That is, the difference is 10^{-4} of the CCT, less than a packet time, showing the tightness of the lower bound.

Dynamics. Fig. 15 illustrates how permutation collectives in fat trees exhibit interesting interactions between data and ACK packets, as each host is at the same time a sender and a receiver. When computing the CCT, we find that there are in fact three successive modes of communication: (1) Data packets only, then (2) interleaved data and ACK packets, and finally (3) ACK packets only. In order to provide a lower bound for the CCT, we need to take these three modes into account.

Assumptions. In the remainder, we focus on a node that receives a flow of data packets going through $H = 6$ hops,

and sends another flow of ACK packets traversing H hops as well. Let T_d, T_a, T_g and T_p respectively denote the data and ACK packet serialization times, the gap time between packets, and the one-way propagation time. Also, let $T'_d = T_d + T_g$ and $T'_a = T_a + T_g$.

Mode 1: only data. The last bit of the first data packet of each message goes through at least a propagation time of T_p and H serialization delays of T_d, thus packet 1 arrives at time $t_1 = T_p + H \cdot T_d$.

When it gets to the receiver, this receiver is in the middle of sending data packet indexed i_1. It needs $T_d + (i_1 - 1) \cdot T'_d$ to send these i_1 data packets. Therefore, putting both equations together yields i_1:

$$i_1 = \left\lceil \frac{T_p + (H-1) \cdot T_d}{T'_d} \right\rceil + 1 \tag{8}$$

Example. With the paper's settings, $i_1 = 78$ data packets. In other words, each sender sends at least 78 consecutive data packets before it sends ACKs. Of course, if $m < i_1$, then it only sends m data packets.

Mode 2: interleaved data and ACK packets. Once the host starts interleaving data and ACK packets, we assume that the host policy performs round-robin between the two, as long as there is data left.

Let's study when the first ACK arrives. Before, we saw that i_1 data packets are sent. We also saw that packet 1 arrives at time $t_1 = T_p + H \cdot T_d$. Therefore, packet i_1 arrives at time $t_{i_1} = t_1 + (i_1 - 1)T'_d$. The first ACK for packet 1 arrives just after, at $t_1^{\text{ACK}} = t_{i_1} + T'_a$.

Afterwards, all packets are sent in an interleaved way, data-ACK-data-ACK-etc. Therefore, the ACK for packet i arrives at $t_i^{\text{ACK}} = t_1^{\text{ACK}} + (i-1)(T'_d + T'_a)$.

Mode 3: only ACKs. At the end, the hosts send the mth packet, and don't need to interleave data and ACKs anymore. This is a symmetric case from i_1: it happens as they send ACK i_2, with $i_2 = m - i_1 + 1$. For example, with $m = 256$ and $i_1 = 78$, it happens at the ACK for packet 179.

Beyond index i_2, we have two constraints. First, ACKs should have time to be transmitted. We get a first lower-bound time $t_i^{\text{ACK}}(1) = t_{i_2}^{\text{ACK}} + (i - i_2)(T'_a)$. Second, ACK packets can obviously only be sent if they acknowledge a data packet that has arrived. Just like packet i_1 arrived just before ACK 1, so does packet i arrive just before ACK $i - (i_1 - 1)$. Thus, we obtain a second lower-bound time: $t_i^{\text{ACK}}(2) = t_{i-(i_1-1)}^{\text{ACK}} + (H-1)(T'_a) + T_p$. Finally, the lower bound is the maximum of both constraints: for $i_2 + 1 \le i \le m$,$t_i^{\text{ACK}} = \max\left(t_i^{\text{ACK}}(1), t_i^{\text{ACK}}(2)\right)$.

B Proof of Thm. 1: Square-Root Scaling

Proof overview. We consider an edge uplink and an aggregation-to-edge downlink. In both cases, the number of packets that reach the queue is a binomial variable, i.e., a sum of independent coin flips, whose mean equals the service. Its excess over the service, and therefore the queue, grows as $\sqrt{m}$. We also prove that no queue grows faster than $\sqrt{m}$, since every link has load at most 1.

Proof of Thm. 1. **No internal traffic.** Since the permutation is a uniform random derangement, the flow destination from any sending host h is uniform over the $n-1$ other hosts, and it lies under h's own edge switch with probability $\frac{w-1}{n-1}$. A union bound over the w hosts of an edge switch shows that it has no internal traffic with probability

$$P_{\text{northbound}} \;\ge\; 1 - w \cdot \frac{w-1}{n-1} \;=\; 1 - \frac{w(w-1)}{2w^3 - 1} \;\ge\; \frac{13}{15} \tag{9}$$

for any fat tree with $w \ge 2$, since the middle term is largest at $w = 2$.

Binomial counts. Fix an index $\ell \in \{1, \dots, w\}$. First, consider uplink ℓ of an edge switch E, and let A_t be the number of packets it receives among the first t packets of each host. If E has no internal traffic, its w hosts send all their packets northbound, and each packet independently takes uplink ℓ with probability $1/w$.

Second, consider the downlink from the ℓ-th aggregation switch of a pod to an edge switch F of this pod. A packet that leaves its source edge switch through uplink ℓ reaches the ℓ-th aggregation switch of the destination pod, so the downlink receives exactly the packets that the w sources sending to hosts under F send through uplink ℓ. Let A_t count them among the first t packets of each source. If F has no internal traffic, each of these packets again takes uplink ℓ independently with probability $1/w$.

In both cases, A_t is binomial with wt trials and parameter $1/w$, with mean t and variance $V = t(1 - 1/w) \ge t/2$. For a binomial variable with N trials and parameter $p \in [1/N, 1 - 1/N]$, the mean absolute deviation is at least $1/\sqrt{2}$ times the standard deviation [7]. Thus $E\left(|A_t - t|\right) \ge \sqrt{V/2} \ge \sqrt{t}/2$, and since $E[A_t] = t$, $E[(A_t - t)^+] = \frac{1}{2}E\left(|A_t - t|\right) \ge \sqrt{t}/4$.

Queues. The queue serves at most one packet per packet duration, so after t packet durations it holds at least $(A_t - t)^+$ packets, provided that the A_t packets have arrived. This is exact for the edge uplink, which receives its packets directly from the hosts; for the downlink, it neglects upstream queueing delays. By Eq. (9), the expected queue after t packet durations is thus at least $\frac{13}{15} \cdot \frac{\sqrt{t}}{4} \ge \frac{\sqrt{t}}{5}$, and its time average over the m packet durations is at least $\frac{1}{5m}\sum_{t=1}^{m}\sqrt{t} \ge \frac{1}{5m}\int_0^m \sqrt{x}\,dx = \frac{2}{15}\sqrt{m}$. Edge uplinks and aggregation-to-edge downlinks each form one fifth of the $5n$ switch queues, leading to the $\sqrt{m}$ lower bound.

Upper bound. Fix a queue X, and let S_b count the packets routed through X among the first b packets of each source. Given the permutation, each packet is routed through X independently, and in each packet duration the expected number of such packets is at most one, e.g., w sources

with probability $1/w$ for an aggregation-to-edge downlink. Thus $S_b - S_a \leq (b-a) + M_b - M_a$, where M_b is a sum of b independent centered terms of variance at most 1, so $E[\max_{b \leq m} |M_b|] \leq 2\sqrt{m}$ by Doob's maximal inequality [28]. Neglecting upstream queueing delays, packets reach X after a fixed time, so by Lindley's formula the queue of X is at most $\max_{a<b}(S_b - S_a - (b-a)) + 1 \leq 2\max_{b \leq m} |M_b| + 1$. Hence $q_{\uparrow}(m) = O(\sqrt{m})$ and $q_{\downarrow}(m) = O(\sqrt{m})$. □

C Proof of Thm. 2: Synchronization

Intuition for synchronization. We prove the synchronization properties of SIMPLE RR and JSQ. The synchronization derives from three properties:

1. Perfect match for first packets. Consider an edge switch without internal traffic. Assume that flow 1 sends the first packet, which is forwarded to uplink port i. Each of the other flows then sends one packet that is forwarded to a distinct uplink port: (i) In RR, because of the round-robin rotation; and (ii) in AR, because the queues of these ports are initially empty, then any port that is assigned a packet gets a queue size of 1 and is removed from consideration for a duration T. At the end, there is a perfect 1-to-1 match of packets to uplink ports.

2. Sticky flows. When the second packet of flow 1 arrives after a period T: (i) in *RR*, it will be sent again to port i due to the round-robin periodic rotation, and (ii) in *AR*, the first packet of flow 1 will have just left, since it takes $L_{\text{PKT}}/C < T$ to service it. Thus port i is the only empty port and the second picks it. In both cases, port i is a sticky port for flow 1. Likewise, each other flow will have its own sticky port.

3. Aggregation switch layout. Now we understand how flows can be stuck onto specific edge-aggregation links. In addition, in a fat tree, a flow stuck to aggregation switch i of a source pod will also be stuck to aggregation switch i of a destination pod (black circles in Fig. 2). Thus, two flows from different source pods but destined to the same edge switch may have picked the same aggregation switch i, leading to contention at the link from the destination aggregation switch to the destination edge switch and to linear growth of its queue.

Proof outline. We first bound the probability that a flow is sticky (Step 1). We then show that an aggregation-to-edge downlink shared by several sticky flows sees its queue grow linearly (Step 2). Since sticky flows from distinct source edge switches pick their aggregation index uniformly and independently, a constant fraction of these downlinks receive at least two sticky flows (Step 3), which yields $q(m) = \Theta(m)$ (Step 4). Finally, we show that the uplink queues stay bounded (Step 5).

For simplicity, the proof assumes that each flow starts with a random jitter uniformly distributed over a packet sending time.

Proof of Thm. 2. Throughout, sources are paced at a fixed rate and are not slowed by congestion control (§2.2); and buffers are assumed infinite. Recall that $w = k/2$ is the number of aggregation-to-edge downlinks entering an edge switch.

Step 1: stickiness probability. We lower-bound the probability p that a flow is *sticky*, i.e., sends all m of its packets through a single aggregation index. We count a flow as sticky only when it satisfies the properties below (the second one only for JSQ), which can only underestimate the queue.

Property 1: no internal traffic. The flow's source edge switch has no internal traffic, i.e., all its w flows are northbound. The number of northbound flows then equals the number of northbound ports, so the first round of packets forms a perfect matching of flows to ports. This happens with probability $P_{\text{northbound}} \geq 13/15$ by Eq. (9).

Property 2: safe flows. For JSQ, stickiness also requires that each packet of the flow finds exactly one port with the strictly shortest queue, namely the port used by its previous packet. Let $T_{SERVICE} = T_{DATA} + T_{IPG}$ be the period, where T_{DATA} is a packet's transmission time and T_{IPG} the inter-packet gap. We call the flow *safe* if no other flow of its edge switch completes a packet transmission within T_{IPG} before or after the flow's own packet completes. Otherwise, the competitor's queue empties either just before the competitor's next packet arrives, which may then pick our flow's port among two empty ones, or just before our flow's next packet arrives, which may then pick the competitor's port. When the flow is safe, its next packet finds a single empty queue: its own previous port, and one packet in each other northbound port. Thus, JSQ then picks the same port deterministically. Since all flows are periodic, the same holds for every subsequent packet. Each of the $w-1$ other flows, whose jitter is independent and uniform over the period, completes a packet within this window of width $2T_{IPG}$ with probability $2T_{IPG}/T_{SERVICE}$, so

$$P_{\text{safe}} = \left(1 - \frac{2T_{IPG}}{T_{SERVICE}}\right)^{w-1}. \tag{10}$$

Stickiness probability. The two properties are independent, since the first depends only on the permutation and the second only on the hosts' jitters. Hence $p = P_{\text{northbound}} \cdot P_{\text{safe}}$ for JSQ. For SIMPLE RR, the round-robin rotation does not depend on queue occupancies, so safety is automatic, $P_{\text{safe}} = 1$, and $p = P_{\text{northbound}}$. In both cases, the index at which a flow sticks is uniform over the w aggregation switches, since it is set by the random initial pointer or tie-breaking together with the hosts' jitters.

Step 2: queue on one downlink. If c sticky flows share an aggregation-to-edge downlink ℓ, each emits one packet per period T, so by the end of period j the link has received at least $c(j-1)$ packets while serving at most j packets, yielding $Q_\ell(j) \geq (c-1)j - c$. Averaging over $j = 1, \ldots, m$, the time-average of ℓ's queue is at least $(c-1)m/2 - O(1)$; the drain phase after the sources stop only increases it. This matters only if a *constant fraction* of the queues carry $c \geq 2$, since one overloaded link among $\Theta(n)$ queues contributes only $\Theta(m/n)$ to the average.

Step 3: a constant fraction of downlinks is overloaded. Let Y_ℓ be the number of sticky flows on downlink ℓ, and fix a destination edge switch E_D. Under a permutation exactly w flows are destined to hosts under E_D, one per host. Their w sources share an edge switch with probability at most $\binom{w}{2}\frac{w-1}{n-w} < \frac{1}{4}$ by a union bound over pairs, so they sit at w distinct source edge switches with probability at least $3/4$ for every $w \geq 2$ ($k \geq 4$).

Let's condition on that event. Distinct source switches then decide independently, so the number B of these flows that are sticky is $\mathrm{Bin}(w,p)$, and each sticky flow falls on one of the w downlinks into E_D uniformly and independently, its index being fixed by its own switch's jitter and initial state. Let O denote the number of occupied downlinks into E_D, so that $\sum_{\ell\to E_D}(Y_\ell-1)^+ = B-O$. Each of the w flows independently ends up as a sticky flow on a given downlink ℓ with probability p/w, so $Y_\ell \sim \mathrm{Bin}(w,p/w)$ and $E[O] = \sum_{\ell\to E_D} P(Y_\ell \geq 1) = w\left(1-\left(1-\frac{p}{w}\right)^w\right)$. Since $E[B] = pw$,

$$E\Big[\sum_{\ell\to E_D}(Y_\ell-1)^+\Big] = E[B]-E[O] = pw - w\left(1-\left(1-\tfrac{p}{w}\right)^w\right) = w\left\{\left(1-\frac{p}{w}\right)^w-(1-p)\right\}. \tag{11}$$

By Bernoulli's inequality $(1-p/w)^w > 1-p$ strictly whenever $w \geq 2$ and $p \in (0,1)$, so the bracket is a positive constant depending only on the topology; it tends to $p-1+e^{-p}$ as w grows. Eq. (11) holds conditionally on the w sources sitting at distinct edge switches; since that event has probability at least $3/4$ and the sum is nonnegative, the unconditional expectation is at least $\frac{3}{4}$ times it. Summing over the $k^2/2$ destination edge switches and using $\frac{k^2}{2}\cdot w = \frac{k^3}{4} = n$,

$$E\Big[\sum_\ell (Y_\ell-1)^+\Big] \;\geq\; \frac{3}{4}n\left\{\left(1-\frac{p}{w}\right)^w-(1-p)\right\} \;=\; \Omega(n). \tag{12}$$

Step 4: combine. A three-level fat tree has $k^2/2$ edge, $k^2/2$ aggregation and $k^2/4$ core switches of radix k, hence $\frac{5k^3}{4} = 5n$ output queues in total. By Steps 2 and 3,

$$\begin{aligned} E\left[q(m)\right] \;&\geq\; \frac{1}{5n}\cdot\frac{m}{2}\cdot E\Big[\sum_\ell (Y_\ell-1)^+\Big] \\ &\geq\; \frac{3}{40}\left\{\left(1-\frac{p}{w}\right)^w-(1-p)\right\}m \;=\; \Omega(m). \end{aligned} \tag{13}$$

For the matching upper bound, at most nm packets ever enter the network, so $q(m)$ is at most linear in m. Hence $q(m) = \Theta(m)$. Since this bound only counts aggregation-to-edge downlink queues, it also gives $q_\downarrow(m) = \Theta(m)$.

Step 5: bounded uplinks. We show by induction over arrival times that no packet ever waits at an uplink. Consider a northbound packet arriving at a switch at time t. Each of the w southern input links of the switch carries at most one packet per period: at an edge switch, because each host sends one packet per period; at an aggregation switch, because by induction the edge uplinks do not queue, and therefore forward their packets as spaced as they receive them. Hence at most $w-1$ other northbound packets arrived in $(t-T,t]$, and those that arrived earlier have already left, since a packet takes less than T to transmit. With SIMPLE RR, these at most w packets are consecutive northbound packets, so they are sent to distinct uplinks. With JSQ, at most $w-1$ uplinks are busy, so the packet joins an empty one. In both cases, the packet is transmitted upon arrival, so every uplink queue holds at most one packet, and $q_\uparrow(m) = O(1)$. Conversely, $q_\uparrow(m) = \Omega(1)$, since by Eq. (9) a constant fraction of the edge uplinks transmit one packet per period. □

D DRB Bounded Queues

Proof of Thm. 3. We want to model the queueing behavior of DRB when using a random permutation. First, at the edge switch uplinks, we saw that we get a sum of periodic flows with random phases. However, since the permutation is random, not all of the flows are necessarily northbound. Thus, the load ρ may be arbitrary. Instead, we model it using its expected value. Any of the w source hosts under an edge switch sends to another host internal to the same edge switch with probability $(w-1)/(n-1)$, therefore the average number of northbound flows is

$$w\cdot\left(1-\frac{w-1}{n-1}\right) \tag{14}$$

Each such flow has a period of wT, where T is the packet transmission time (including inter-frame gap). We model this system as $N\cdot D/D/1$, where N is this expected number of northbound flows, each of period $D = wT$ at a given uplink port, and the load is $\rho = N/w < 1$. We model the superposition of periodic arrivals as a Gaussian process as k increases [2], together with a negative drift since $\rho < 1$. Since the queue size cannot be negative, we use a truncated normal distribution. This queue is stationary: its distribution depends on the number of flows and on ρ, but not on how long the flows keep sending, so the modeled average queue is independent of m, i.e., $\Theta(1)$.

In the same way, for aggregation-to-core links, we neglect intra-pod traffic, and get the following expected number of inter-pod flows per source pod:

$$w^2\cdot\left(1-\frac{w^2-1}{n-1}\right) \tag{15}$$

We then apply the same model. Note that the fit is not exact anymore, as traffic coming from the edge switches loses its periodicity and therefore it is important to emphasize that *the ND/D/1 model is only an approximation*. Downlinks are an even coarser approximation, since each layer makes traffic less periodic. □

An exact bound at the edge uplinks. For the first queueing layer, the model can be replaced by a deterministic bound.

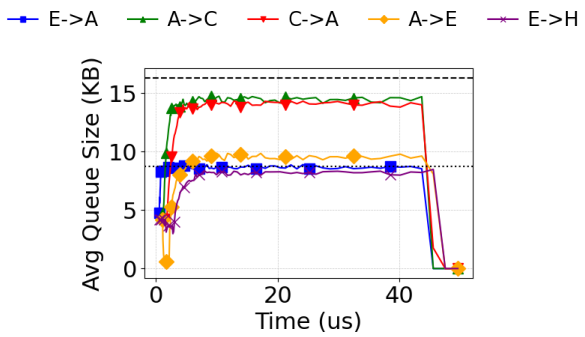

Figure 16: Queue models for DRB. The average queue size at each of the 5 switch layers is compared against queueing models (dotted and dashed horizontal lines).

As above, each northbound flow appears at a given uplink port with period wT, so it contributes at most $\frac{t-s}{wT}+1$ packets to that port in any interval $(s,t]$. With at most w northbound flows, the port receives at most $\frac{t-s}{T}+w$ packets in $(s,t]$ while serving one packet per T, so its queue never exceeds $w+1$ packets, whatever the message size m. The other layers lose this periodicity, which is why we resort to the $ND/D/1$ approximation there.

Validation. Fig. 16 uses $n = 1,024$ hosts with a 4 MB message to show the behavior of DRB at scale. It shows the average queue sizes at each of the 5 switch queueing layers: (i) edge to aggregation ($E \to A$), (ii) aggregation to core ($A \to C$), (iii) core to aggregation ($C \to A$), (iv) aggregation to edge ($A \to E$) and (v) edge to host ($E \to H$). It illustrates that the approximation is very close for the first E→A layer. The model then is a bit pessimistic for the two core layers. It has three main sources of errors: (i) the traffic periodicity is broken by queueing at successive switches, (ii) the load varies and is different at each pod, and (iii) we only use a Gaussian approximation. Finally, the two downlink edge layers are relatively close to the model as well despite the loss of periodicity.

E ETHEREAL Bounded Queues

Proof of Thm. 4. We use the same model as for DRB (Appendix D), with a single change. Under a permutation, each host sends a single flow in DRB, which ETHEREAL [1] splits into s subflows of m/s packets, one per path to the destination. Since the NIC serves its queue pairs in round-robin [1], each subflow is a periodic source that sends one packet every sT, with a random phase. Hence the only change is that we model subflows, rather than flows, as sources: more sources, each with a longer period, for the same total rate.

This change does not affect the boundedness of the queue. First, each path receives exactly m/s packets of each flow, so, as with DRB, all downlinks into a switch carry the same load and every link has load $\rho \leq 1$. Second, each queue is still an $N \cdot D/D/1$ queue, which is stationary, so its average does not depend on m, i.e., $q(m) = \Theta(1)$. □

An exact bound at the edge uplinks. As for DRB, the first layer admits a deterministic bound. Consider an edge uplink and a northbound flow of its edge switch. If the flow is inter-pod, w of its $k^2/4$ subflows use that uplink, each sending one packet every $\frac{k^2}{4}T$; if it is intra-pod, one of its w subflows does, sending one packet every wT. Either way, the flow contributes at most $\frac{t-t'}{wT}+w$ packets to the uplink in any interval $(t',t]$. With at most w northbound flows, the uplink receives at most $\frac{t-t'}{T}+\frac{k^2}{4}$ packets while serving one per T, so its queue never exceeds $k^2/4+1$ packets, whatever m. The constant is larger than DRB's $w+1$ because the random order of the queue pairs does not alternate between aggregation switches, so a flow may send several consecutive packets through the same uplink.

F Proofs and Details of OFAN's Properties

F.1 Correctness of Consolidation

Notation. Throughout this appendix, H denotes a host, E an edge switch, A an aggregation switch and C a core switch, with subscript S for the source side and D for the destination side. Thus E_S and A_S are a source edge switch and a source aggregation switch, and H_D, E_D and A_D are a destination host, its edge switch, and an aggregation switch of its pod. A *downlink* is a southbound link, i.e., $C \to A_D$, $A_D \to E_D$ or $E_D \to H_D$.

Load-balancing invariant. We start by formally formulating the goal of a switch DR in general (and OFAN in particular) in a network without link failures. It is clear that we want a perfect LB locally among uplink ports. But we also want to make sure that any source edge switch E_S load-balances traffic as equally as possible across the links in the destination pods. We want to formulate a target ideal invariant that this load balancing should satisfy. However, we need to be careful in the formulation. The traffic that E_S load-balances depends on the destinations of the flows that cross E_S. These destinations depend on the arbitrary permutation pattern, and are therefore not uniform. For example, E_S cannot send the same amount of traffic to all edge switches in the network. Instead, we condition the invariant on the destination switch.

Invariant 1 (Load-balancing invariant). *Given any source switch X in a source pod and any destination switch Y in a destination pod, X should send traffic at the same rate through all of the southbound links to Y that are reachable from X.*

The invariant implies a perfect load balancing at downlinks:

Property 1. *If Inv. 1 holds for source aggregation switches (X) and destination aggregation switches (Y), and likewise*

for source edge switches and destination edge switches, then all southbound links to any switch have equal load.

Proof. Partition the traffic on the downlinks into a switch by its source at the layer above: the load on each core-to-A_D link is the sum of the rates of the source aggregation switches that reach it, and the load on each A_D-to-E_D link is the sum of the rates of the source edge switches that reach it. If each source sends at the same rate on all the links it reaches, these sums are equal for all the downlinks into a given switch. □

Consolidation. To achieve the above target invariant, we use *consolidation into virtual flows*. The key idea is that each source switch consolidates into a single shared virtual flow all of its northbound packets destined to a shared destination Y. Then, it load balances these packets across all of its uplink ports. However, we still need to carefully define the level at which to consolidate flows.

Mandatory waypoints. Consider a source aggregation switch A_S. A destination host H_D that is reachable through its uplinks is necessarily in a different pod. Moreover, a property of the three-level fat tree is that when sending traffic from A_S to H_D, any packet *must* reach a given destination edge E_D and a given destination aggregation switch A_D in the destination pod, and cannot reach other edge or aggregation switches. For example, the third aggregation switch of the source pod can only send traffic destined to H_D through the third aggregation switch A_D of the destination pod, then through the edge switch E_D that is connected to H_D. We can think of A_D and E_D as *mandatory waypoints*. They stand in contrast with the many possible core switch waypoints that could be used to reach H_D.

As a consequence, all of the traffic going through A_S and destined to H_D could be consolidated into a single virtual flow using several consolidation levels: (i) per destination host H_D, (ii) per destination edge E_D, and (iii) per destination aggregation A_D. The following result shows that they can all provide a first performance guarantee.

Theorem 8. *Source aggregation switch A_S satisfies Inv. 1 whether it consolidates traffic per destination host, edge or aggregation switch.*

Proof. Whatever the consolidation level, for any destination belonging to that level in the destination pod, A_S always sends traffic at the same rate to all the cores it is connected to. Thus A_D, the aggregation switch of the same index in the destination pod, which is connected to those same cores, also receives the same rate from each of them, and therefore satisfies the invariant. In addition, the only downlinks A_S can reach beyond A_D are $A_D \to E_D$ and $E_D \to H_D$, and it can reach no other downlink into E_D or into H_D. Being the only ones, they trivially satisfy the invariant. □

Consider now a source edge switch E_S. Its mandatory waypoints for a given destination host H_D are only H_D and E_D. Since it can reach H_D using several aggregation switches, none is mandatory. Thus, E_S can consolidate traffic (1) per destination host H_D and (2) per destination edge E_D. We find that both work out:

Theorem 9. *No matter the consolidation level, if both the source edge switch E_S and the source aggregation switch A_S consolidate traffic, then A_S satisfies Inv. 1 for the downlinks into destination aggregation switches, and E_S for the downlinks into destination edge switches and hosts. Hence all the downlinks into any switch have equal load.*

Proof. We structure the proof based on the destination type.

(1) For any destination aggregation switch A_D, the source edge switch E_S has a single source aggregation waypoint A_S: if the destination is in the same pod then $A_S = A_D$, and otherwise A_S is the aggregation switch of the same index in the source pod. All of E_S's traffic for A_D is therefore sent to A_S. Since A_S shares its pointer with the other edge switches of its pod, E_S's own share need not be spread evenly over the cores; but the aggregate that A_S sends satisfies the invariant for A_D by Thm. 8, which is all that Property 1 needs for the core-to-A_D links.

(2) For any destination edge switch E_D, the source edge switch E_S splits its traffic for E_D equally across all the aggregation switches of its own pod, sending the same rate R to each. In turn, the rate leaving any source aggregation switch A_S is carried to the aggregation switch A_D of the same index in the destination pod (with $A_D = A_S$ when the two pods coincide). Hence E_D receives the same rate R on each of its downlinks from the aggregation switches of its pod.

(3) Finally, destination host H_D is only connected through a single link to the network, so the invariant holds trivially. □

Since it is clear that pointers also load-balance perfectly locally among uplink ports, the above theorem means that LB achieves its dual goals both at uplinks and at downlinks.

F.2 Low-Cost Per-Destination Port Orders

Requirements. Ofan keeps, at each switch, one round-robin pointer per destination (per destination edge switch at an edge switch, and per destination pod at an aggregation switch) and per packet type (e.g., data or ACK). Each pointer is an index $i \in \{0, \ldots, w-1\}$ into an order of the switch's w uplinks, starting at a random position. Two requirements drive the choice of order. First, starting from any position, a pointer must visit every uplink exactly once per w packets, so traffic is split evenly. Second, pointers that are active at the same time at the same switch should follow different orders; otherwise they tend to select the same uplink at the same time. Storing an independent random permutation per destination meets both requirements, but costs $\Theta(w \log w)$ bits per destination.

OFAN operation. Instead, OFAN computes the order from a single global permutation P of $\{0, \ldots, w-1\}$, shared by all

switches, and a per-pointer identifier

$$id = d + N_d \cdot t, \tag{16}$$

where d is the destination (edge switch or pod), N_d is the number of possible destinations, and $t \in \{0,1\}$ distinguishes data from acknowledgments. We split this identifier into its low-order and high-order parts,

$$id_1 = id \bmod 2^b, \qquad id_2 = \lfloor id/2^b \rfloor, \tag{17}$$

where $b = \lceil \log_2 w \rceil$. The i-th port of the order is then

$$\pi(i) = \left(\mathrm{xor}_w(P[i], id_1) + id_2\right) \bmod w, \tag{18}$$

where $\mathrm{xor}_w(v,x)$ equals $v \oplus x$ if $v \oplus x < w$ and v otherwise.

Correctness. Each step maps $\{0,\ldots,w-1\}$ onto itself one-to-one. Adding a constant modulo w clearly does. For xor_w, since XOR is its own inverse, if v is kept because $v \oplus x \geq w$, no other value $u < w$ can satisfy $u \oplus x = v$, since that would require $u = v \oplus x \geq w$. Hence π is a permutation, and every pointer visits each uplink exactly once per w packets from any starting position. Under link failures, ports that no longer lead to the destination are skipped, so each remaining valid port still receives exactly one packet per round.

Diversity from two non-collapsing operations. Using a single kind of operation yields few orders: all additions combine into one shift, giving at most w orders, and all XORs combine into one XOR, giving at most 2^b. XOR and addition, however, do not reduce to one another: addition preserves arithmetic differences and breaks bitwise ones through carries, while XOR does the opposite. Driving the XOR with id_1 (the low-order bits of id) and the addition with id_2 (its high-order bits) therefore produces up to $2^b \cdot w$ combinations. For $w = 6$ they yield 30 distinct cyclic orders, versus 6 with additions alone, so the few destinations served concurrently by a switch almost always follow different orders.

Computational cost. The per-packet work is one lookup of $P[i]$, one XOR, one comparison, one addition, and one conditional subtraction, all on $\lceil \log_2 w \rceil$-bit values. The constants id_1 and id_2 mod w can be derived from the destination on the fly.

Storage cost. Beyond the pointers themselves, a switch stores nothing per destination. P is a single global constant of $w\lceil \log_2 w \rceil$ bits that can be hard-wired and is the same for all switches. P is drawn once offline and chosen so that its w cyclic shifts define w distinct orders.

F.3 Scaling Impact vs. DRB

Each switch can consolidate at several levels (Thms. 8 and 9), and OFAN always chooses the highest one: the higher the consolidation level, the fewer pointers a switch holds, and the smoother its load balancing, since each pointer coordinates more flows. We now quantify this gain vs. DRB.

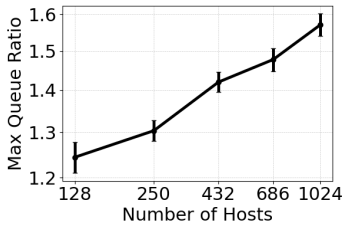

Figure 17: Ratio of the maximum queue size of DRB to that of OFAN as a function of the number of hosts.

Scaling impact on pointer numbers. Consider a random permutation with inter-pod traffic only. In OFAN, aggregation switches consolidate traffic per destination aggregation switch. Thus, a destination aggregation switch can only receive up to $k-1$ different virtual flows, since only the $k-1$ source aggregation switches of the same index, one per other pod, can reach it. However, in DRB, they virtually consolidate traffic per destination host, hence there are $k^2/4$ flows, corresponding to the $k^2/4$ hosts in the destination pod. Thus, the consolidation gain in the number of pointers is at least $k^2/(4(k-1)) > k/4$, increasing with the network size.

Scaling impact on queue sizes. Under the same traffic, Fig. 17 plots the ratio of the maximum queue size of DRB to that of OFAN as a function of network size, with 30 runs per size. The ratio increases with the network size, i.e., OFAN outperforms even more as we scale the network.

F.4 Model with Bounded Average Queue Sizes

We now build a model for the queue size under OFAN.

Proof of Thm. 5. The model follows similar lines for OFAN as for DRB. However, since we consolidate pointers, we evaluate the number of distinct pointers, and model as if all pointers carried the same number of periodic flows, ignoring the variability across pointers.

We start with the edge model. We want to determine the expected number of pointers at the edge switch. A single source edge switch comprises w flows (balls), and these flows can be destined to $k^2/2$ edge switches (bins). The expected number of distinct destination edge switches is:

$$\frac{k^2}{2} \cdot \left(1 - \left(1 - \frac{1}{k^2/2}\right)^w\right).$$

Using Taylor expansion, we find that this is about $w - 1/4$.

Likewise, at each aggregation switch, the number of northbound pointers is at most $k-1$, as there are $k-1$ other pods and therefore other aggregation switches that can be destinations. The expected number of pointers also converges to $k-1$ as k grows. Thus, we obtain a new model, which neglects many aspects of the algorithm that can be too complex to capture, and in particular the variability in the periods of

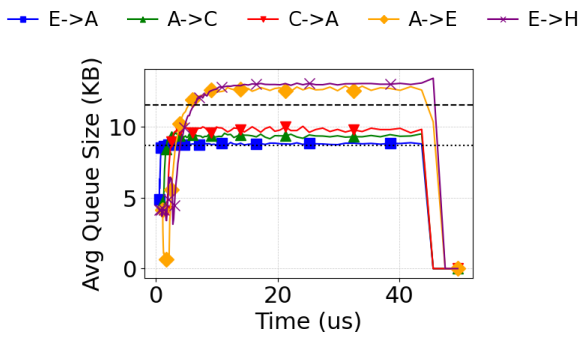

Figure 18: Queue models for OFAN. The average queue size at each of the 5 switch layers is compared against queueing models (dotted and dashed horizontal lines).

each virtual flow. Again, *the ND/D/1 model is only an approximation*. As for DRB, the modeled queue is stationary and therefore independent of m. □

An exact bound at the edge uplinks. As for DRB (Appendix D), the first queueing layer admits a deterministic bound. A pointer shared by f flows advances f times per period T, and therefore visits each uplink port exactly f times per w periods, i.e., it contributes at most $f\left(\frac{t-s}{wT}+1\right)$ packets to a given port in any interval $(s,t]$. Since the (at most w) northbound flows of an edge switch are split among its pointers, the port receives at most $\frac{t-s}{T}+w$ packets in $(s,t]$ while serving one packet per T, so its queue never exceeds $w+1$ packets, whatever the message size m.

Validation. Fig. 18 illustrates the model for OFAN under the same conditions. We can see that the edge model holds well (dotted line). However, as expected, the models beyond have a slightly poorer fit. In particular, the downlink queues are harder to model exactly.

F.5 Dynamic First Pointer

In a network without link failures, we find that if (i) the first pointer is not picked at random, but rather chooses the shortest queue at the arrival of the first packet, then (ii) it cycles in an unshuffled round-robin order, then CCT is slightly improved. However, switches would need to know the start of a collective, or restart pointers after a timeout; we do not further evaluate this variant in this paper.

G Proof of Thm. 6: Optimal Packet Size

Proof of Thm. 6. The CCT can be modeled as the sum of (i) the time to send the message, (ii) the time to go through the queues, and (iii) the zero-load RTT. Only the first two components depend on the packet size P, so let's focus on them.

(i) The time to send the message is

$$\frac{D}{P-H}\cdot\frac{P}{C}, \tag{19}$$

since we need $\frac{D}{P-H}$ packets to send a message of D bytes using $P-H$ payload per packet, and each packet of size P that is sent at rate C takes time $\frac{P}{C}$.

(ii) The time to go through the queues is

$$\alpha\cdot\frac{P}{C}, \tag{20}$$

by assumption, since the queues hold α packets of size P.

Adding the two components and dropping the common factor $\frac{1}{C}$, we want to minimize

$$f(P) := P\cdot\left(\frac{D}{P-H}+\alpha\right) = \frac{PD}{P-H}+\alpha P. \tag{21}$$

Differentiating with respect to P, we get

$$f'(P) = -\frac{D\cdot H}{(P-H)^2}+\alpha. \tag{22}$$

Zeroing the derivative gives $(P-H)^2 = \frac{DH}{\alpha}$, that is $P-H = \sqrt{\frac{H}{\alpha}D}$, which is the claimed result. It is indeed a minimum, since $f''(P) = \frac{2DH}{(P-H)^3} > 0$ for every $P > H$. □

Note that Fig. 8 assumes as a first approximation $H = 82$ B, accounting for the header and inter-frame gap components, and $\alpha = 10$ data packets.

H Proof of POFAN Queueing Model (Thm. 7)

Proof of Thm. 7. We use the same model as in the proof of Thm. 5 (Appendix F), with a single change. In OFAN, each key ((i) destination edge switch at an edge switch or destination pod at an aggregation switch, and (ii) size class) has a single pointer per switch; in POFAN, each key has one independent pointer per ingress pipe, and a packet uses the pointer of the pipe it enters through. Each of these pointers still rotates over all the switch's uplinks. Hence the only change is that more pointers, i.e., more periodic sources, feed each queue: up to P times more for P ingress pipes.

This change does not affect the boundedness of the queue, for two reasons. First, the load on every link stays balanced: each pipe pointer spreads its own traffic equally over its uplinks, so the sum of the pipe pointers of a key does too, exactly like OFAN's single pointer. The proofs of Thms. 8 and 9 therefore hold unchanged, and every link keeps a load $\rho \leq 1$ by Property 1. Second, the modeled $N\cdot D/D/1$ queue is stationary for any number N of sources: a larger N changes the value of the average queue, but not the fact that it is independent of m. Since N is at most P times OFAN's number of pointers, which does not depend on m, we get $q(m) = \Theta(1)$. □

An exact bound at the edge uplinks. The deterministic bound of Appendix F also carries over unchanged, since it only used that the (at most w) northbound flows of an edge switch are split among pointers such that a pointer shared by f of these flows visits every uplink port f times per w periods, whatever the split. Hence the queue of any edge uplink under POFAN is at most $w+1$ packets, whatever the message size m.

The cost of pipes. The price of the pipes is a larger number of pointers, hence a larger constant. At an edge switch the difference is small, since OFAN's $\approx w-1/4$ edge pointers mostly carry a single flow each already. At an aggregation switch with P ingress pipes, however, OFAN's at most $k-1$ pointers become up to $P(k-1)$, e.g., $P = k/(2W) = 8$ for $k = 64$ and pipes of $W = 4$ ports. This is consistent with the gap between POFAN and OFAN widening with the network size in Fig. 6, while both remain bounded in m. As for OFAN, the model assumes that the pointers are randomly initialized so that they do not lock in phase, and it is only an approximation beyond the edge layer.

I OFAN's Hardware Implementation

Concurrent access challenge. Implementing OFAN's per-destination-group pointers at line rate raises a structural challenge: a switch processes many packets per cycle in parallel, arriving from different input ports, and several of these packets can be destined to the same destination edge or aggregation switch within the same cycle. Such packets must all access the *same* OFAN round-robin (RR) pointer, creating a read-modify-write hazard on a single shared register.

Combining network. This concurrent-access problem is not new: it is a classic problem in parallel computing, where many concurrent processors issuing fetch-and-add requests to a shared counter can create the same kind of contention [46,50]. The classical fix, called a *combining network*, is to merge duplicate concurrent requests together before they reach the shared counter, so the counter is only touched once no matter how many requests asked for it.

That said, these classical papers solve a more general problem, where requests to a huge, dynamically-addressed memory are routed hop-by-hop through a large network (e.g., an omega or butterfly network) and only combine when they happen to collide en route. This general machinery is costly and has variable latency, which does not fit a switch pipeline that needs guaranteed, fixed per-cycle latency. Instead, we want to solve our narrower problem, over a small fixed set of N local pointers, using a much cheaper, fixed-latency, purpose-built circuit, described next.

OFAN pipeline. We can implement OFAN's pointers with a simple, fixed-latency pipeline that processes a batch of w packets in parallel, where $w = k/2$ is the number of downlink ports (only packets arriving from downlinks need a pointer decision, since uplink paths are chosen while downlink paths are unique). The pipeline has four stages: (1) sort the batch by pointer key, so packets sharing a pointer become adjacent; (2) combine the counts of packets sharing each pointer key; (3) update the shared pointer register for each key; and (4) decombine the updated pointer values back to each packet in the batch. We illustrate the pipeline in our fat-tree example of §4 with $k = 64$-port switches, i.e., at most $w = k/2 = 32$ packets arriving in parallel from downstream inputs, 32 uplink ports to cycle through, and $N = 2K$ edge switch destinations.

(1) Sort. First, a w-input bitonic sorting network [17] sorts the batch by pointer key, so packets sharing a pointer become adjacent.

(2) Combine. A combining tree [25] computes, in $\log_2(w)$ stages, how many packets in the batch share each key; the last packet of each group becomes its *master*, carrying the group's match count.

(3) Update. A register file of N per-pointer counters is updated: each master atomically fetches its pointer's current value and advances it, modulo the number of uplink ports, by the group's match count.

(4) Decombine. Finally, a reverse network decombines the fetched value back to every packet in its group (offsetting it by the packet's position within the group), and an unsorting network—the inverse of the initial permutation—restores the original per-lane order, so each packet leaves the pipeline carrying its own assigned uplink port. For our running example, this requires only a 32-input sorting/combining tree and $2K$ registers per edge switch, a practical hardware budget.

Example. Suppose 3 of the w packets in a batch are all destined to the same edge switch, which has $k/2 = 32$ uplink ports, so its pointer takes values in $\{0,1,\ldots,31\}$ and currently holds the value 1. In stage (1), sorting by pointer key brings these 3 packets together, adjacent to each other. In stage (2), the combining tree counts the 3 matches and marks the last of them as the *master*, carrying a match count of 3. In stage (3), this master atomically fetches the pointer's current value (1) and advances it by 3 mod 32, leaving the pointer at $(1+3) \bmod 32 = 4$ for the next batch. In stage (4), the 3 packets are decombined: each is assigned the fetched value offset by its position within the group, i.e., ports 1, 2, and 3 respectively, and the unsorting network routes each result back to the packet's original lane. Thus, the pointer is touched exactly once for all 3 packets, yet each packet still receives its own distinct, correctly-ordered round-robin port, exactly as if the 3 requests had been processed one at a time.

Cost estimate. The pipeline requires N registers for the pointers, plus a w-input sorting network, a w-input combining tree, and a matching decombining/unsorting network of the same size. For our running example, recall each edge switch keeps one pointer per (destination edge switch, packet-size quantization group) pair; with $2K$ destination edge switches and 2 quantization groups (data, ACKs), $N \approx 4K$. Each pointer needs $\log_2(32) = 5$ bits to represent its value (one of the 32

uplink ports), so the register file is only $N \times 5 \approx 20K$ bits ($\approx$ 2.5KB). The rest of the pipeline is dominated by two contributions: (i) the W=32-input sorting and unsorting networks, each with 15 stages of 16 compare-and-swap cells (roughly 80 gates per cell, for a 12-bit pointer-key comparator, wide enough to address $N \approx 4K$ keys, plus payload-swap muxes), for $\approx 20K$ gates apiece; and (ii) the register file's read-modify-write logic, dominated by the $20K$ flip-flops themselves at ≈ 4 gate-equivalents each, i.e., $\approx 82K$ gates. The combining tree adds only a couple thousand more gates ($\approx$ 80 small adder nodes for the segmented match counts). Altogether, this is on the order of 10^5 logic gates ($\approx 127K$). At a 1 GHz clock and a conservative 10% switching activity, this amounts to only a few milliwatts of dynamic power ($\approx$ 6mW).

Both figures are negligible next to the resources of the very chip we used as our pipe-architecture example: the BCM56980 (Tomahawk 3) datasheet lists a typical power draw of 300 W and 64 MB of integrated packet-buffer memory [13], independently corroborated by [83]. So OFAN's pipeline's $\approx$ 6mW would add only $\approx$ 0.002% to such a chip's power budget, and its $\approx 20K$-bit register file is only $\approx$ 0.0037% of that chip's 64MB ($\approx$ 537 million bits) of buffer memory—both entirely negligible.

Other details. Implementing OFAN only requires modifying the core scheduling algorithm in a switch. No change is required at endhosts and OFAN can continue to use existing routing protocols such as BGP W-ECMP [6, 37, 38, 48, 53, 60, 61, 65] to propagate information about failures etc.