

Packet Forwarding with a Locally Bursty Adversary

Will Rosenbaum  

Amherst College, MA, USA

Abstract

We consider packet forwarding in the adversarial queueing theory (AQT) model introduced by Borodin et al. We introduce a refinement of the AQT (ρ, σ) -bounded adversary, which we call a *locally bursty adversary* (LBA) that parameterizes injection patterns jointly by edge utilization and packet origin. For constant $O(1)$ parameters, the LBA model is strictly more permissive than the (ρ, σ) model. For example, there are injection patterns in the LBA model with constant parameters that can only be realized as (ρ, σ) -bounded injection patterns with $\rho + \sigma = \Omega(n)$ (where n is the network size). We show that the LBA model (unlike the (ρ, σ) model) is closed under packet bundling and discretization operations. Thus, the LBA model allows one to reduce the study of general (uniform) capacity networks and inhomogenous packet sizes to unit capacity networks with homogeneous packets.

On the algorithmic side, we focus on information gathering networks – i.e., networks in which all packets share a common destination, and the union of packet routes forms a tree. We show that the Odd-Even Downhill (OED) forwarding protocol described independently by Dobrev et al. and Patt-Shamir and Rosenbaum achieves buffer space usage of $O(\log n)$ against all LBAs with constant parameters. OED is a local protocol, but we show that the upper bound is tight even when compared to centralized protocols. Our lower bound for the LBA model is in contrast to the (ρ, σ) -model, where centralized protocols can achieve worst-case buffer space usage $O(1)$ for $\rho, \sigma = O(1)$, while the $O(\log n)$ upper bound for OED is optimal only for local protocols.

2012 ACM Subject Classification Theory of computation → Routing and network design problems; Networks → Packet scheduling; Theory of computation → Distributed algorithms; Theory of computation → Distributed computing models

Keywords and phrases packet forwarding, packet scheduling, adversarial queueing theory, network calculus, odd-even downhill forwarding, locally bursty adversary, local algorithms

Digital Object Identifier 10.4230/LIPIcs.DISC.2022.34

Acknowledgements This work was born from numerous discussions with Boaz Patt-Shamir, to whom I am eternally grateful. I thank the anonymous reviewers for their thoughtful commentary which helped improve this paper.

1 Introduction

Routing and forwarding are fundamental operations in the study of networks. In this context, commodities – for example, data packets, fluid flows, or physical objects – appear at various places in a network, and must be transferred to prescribed destinations. Movement is restricted by the network’s topology. The goal is to get the commodities from source to destination as efficiently as possible. *Routing* is the process of determining routes for the commodities to follow from source to destination, while *forwarding* determines the particular schedule by which items – which we will henceforth refer to as *packets* – move in the network. In this work, we focus on the process of forwarding packets, assuming their routes are pre-determined.

Two well-studied models of packet forwarding in networks are the adversarial queueing theory (AQT) model introduced by Borodin et al. [2] and the network calculus model introduced by Cruz [4, 5]. In both models, packets are assumed to have prescribed routes from source to destination. Both models also parameterize packet arrivals in terms of long-



© Will Rosenbaum;

licensed under Creative Commons License CC-BY 4.0

36th International Symposium on Distributed Computing (DISC 2022).

Editor: Christian Scheideler; Article No. 34; pp. 34:1–34:18

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

term average rates and short-term burstiness in order to disallow trivially infeasible injection patterns that exceed network capacity constraints. AQT and network calculus also differ in some crucial ways. AQT examines injections of discrete, indivisible packets at discrete time intervals, and forwarding occurs in synchronous rounds. In network calculus, on the other hand, packets are modeled as continuous flows and forwarding is a continuous-time processes. Nonetheless, these flows can be discretized (or “packetized”) to be processed discretely. One of the goals of this paper is to draw tighter connections between analogous parameters in the AQT and network calculus models under the process of discretization.

In both AQT and network calculus, one natural measure of efficiency is the buffer space usage of nodes in the network. That is, how much memory is required at each buffer in order to store packets that are *en route* to their destinations. Traditionally, AQT has focused on a qualitative measure of space usage, called *stability*, which merely requires that the space usage of a protocol remains bounded (by some function of the network parameters) for all time. A notable early exception is the work of Adler and Rosén [1], which gives a quantitative buffer space upper bound for longest-in-system scheduling when the network is a directed acyclic graph.

AQT has also traditionally focused on *greedy* forwarding protocols – i.e., protocols for which every non-empty buffer forwards as many packets as it can in each round (subject to capacity constraints). A more recent series of work [8, 6, 10, 11, 9] initiated by Miller and Patt-Shamir [8] studies quantitative buffer space bounds for non-greedy forwarding policies. In particular, these works show that in restricted network topologies (single-destination paths and trees), non-greedy forwarding protocols can achieve significantly better buffer space usage than greedy protocols. Specifically, non-greedy centralized forwarding protocols can achieve $O(1)$ buffer space usage [8, 9], while $\Theta(\log n)$ buffer space is necessary and sufficient for local (distributed) protocols [6, 10] (where n is the number of buffers in the network). The work of Patt-Shamir and Rosenbaum [11] shows there is a smooth trade-off between a protocol’s *locality* and optimal buffer space usage: if each node determines how many packets to forward based on the state of its distance d neighborhood, then $\Theta(\frac{1}{d} \log n)$ buffer space is necessary and sufficient. These bounds are in contrast to greedy protocols, which require $\Omega(n)$ buffer space in the worst case.

The bounds described in the preceding paragraph refer to the AQT injection model in which edges in the network have uniform unit capacities – only one packet may cross any edge in a given round – and the average injection rate ρ satisfies $\rho \leq 1$, and the burst parameter σ satisfies $\sigma = O(1)$ (cf. Definition 2.1). The algorithms can be generalized to general uniform edge capacities ($C \geq 1$ packets can cross each edge in a round), but the generalized algorithms are both more cumbersome to express, and correspondingly subtle to reason about (see, e.g., Section 1.1 in [6]). When dealing with general capacities, discretizations of continuous flows, and heterogeneous (indivisible) packets, a natural strategy is to bundle packets into “jumbo packets” [12]. This procedure can, however, lead to large bursts in the appearance of jumbo packets in the network, even if the injection process has a small burst parameter (see e.g., Remark 2.10). Thus applying the *analysis* of the relatively simple unit-capacity versions of algorithms in [8, 6, 10, 11, 9] directly to jumbo packets may not show any improvement over greedy algorithms.

1.1 Our Contributions

In this paper, we introduce a refinement of Borodin et al. [2]’s (ρ, σ) parameterization of injection patterns, that we call a the *local burst model* (see Definition 2.2). In addition to the asymptotic rate ρ and *global* burst parameter σ , the local burst model has a third

parameter, β – the *local* burst parameter, that accounts for simultaneous bursts occurring at distinct injection sites. Thus, for small values of β , a locally (ρ, σ, β) -bounded injection pattern may still allow for large (e.g, $\Omega(n)$) simultaneous packet injections, so long as not too many packets are injected into the same buffer.

We prove that locally bursty injection patterns are essentially characterized as discretizations of what we call “locally dependent flows” (Definition 2.6) with similar parameters – see Lemmas 2.7 and 2.9. We use this characterization to show that applying packet bundling to a locally bursty injection pattern yields another locally bursty adversary with similar parameters – see Proposition 3.1. Consequently, any space efficient algorithm for unit capacity networks and homogeneous unit sized packets can be applied as a black-box to bundled jumbo packets to achieve similar buffer space usage to the unit capacity case (Corollary 3.2). We then show that a small modification of the framework can also be applied to the setting of heterogeneous packet sizes.

On the algorithmic side, we analyze the *odd-even downhill* (OED) forwarding protocol of [6, 10] against locally bursty adversaries. We show that for constant ($O(1)$) parameters ρ, σ , and β , OED achieves worst-case buffer space $O(\log n)$ in information gathering networks of size n – see Theorem 4.2. This result is strictly stronger than the analyses of [6, 10], as there are locally bursty injection patterns with $\beta = O(1)$ that can only be realized in the classical (ρ, σ) model for $\sigma = \Omega(n)$. Combining our analysis of the OED protocol together with flow discretization and/or bundling, buffer space of $O(\log n)$ can be achieved for forwarding with general capacities, heterogeneous packets, and continuous flows – see Section 4.1.

Finally, in Section 5, we prove a matching lower bound of $\Omega(\log n)$ for any *centralized randomized* protocol against locally bursty adversaries (Theorems 5.1 and A.1). This lower bound is in contrast to the deterministic upper bounds of [8, 11, 9] which show that $O(1)$ buffer space is achievable for (ρ, σ) -bounded adversaries for centralized and “semi-local” protocols. Thus, our lower bound shows that the local burst model (with constant parameters) gives the adversary strictly more power to inflict large ($\omega(1)$) buffer space usage on *centralized* algorithms. The performance of the asymptotically optimal *local* protocol is the same for the local burst and traditional injection models, and in the local burst model, the (local) OED protocol is asymptotically optimal, even when compared to centralized protocols.

1.2 Discussion of Our Results

The most natural application domains for this work are networks consisting of tightly synchronized nodes, such as network-on-chips (NoCs) [7], software defined networks (SDNs) [13], and sensor networks. In these contexts, trees and grids are common network topologies. (In the case of grids, “single bend” routing allows one to treat the network essentially as a disjoint union of paths.) Thus, while the topologies we consider are highly restricted, the family of topologies is a fundamental and frequently used family for applications in which our techniques might be applied.

In NoCs, SDNs, and sensor networks, the rate and source of packet injections may be highly variable. Thus, the parameterizations of packet injections in both the standard AQT model and the network calculus model may be too coarse to model the actual buffer space requirement of observed injection patterns effectively. In the AQT model, allowing for multiple simultaneous packet injections into different buffers requires a large burst parameter, σ , even though the resulting injection pattern may be handled using buffer space $\ll \sigma$ (see Example 2.4). On the other hand, the traditional network calculus model does not account for dependencies between rates of packet injections into different buffers over time. Thus, a standard analysis may severely overestimate the bandwidth or buffer space required to

handle a given injection pattern. Our locally bursty injection model (and its continuous analogue described in Section 2.2) refine both the AQT and network calculus models so as to give more precise bounds on the buffer space requirement of many natural packet injection patterns.

Together with the upper bounds of [8, 11, 9], Theorems 5.1 and A.1 imply that the locally bursty injection model (with constant parameters) gives an adversary strictly greater power to inflict large buffer space usage against *centralized* and semi-local forwarding protocols. However, Theorem 4.2 implies that the *local* OED forwarding protocol achieves asymptotically optimal buffer space usage. Thus, for locally bursty injection patterns, there is no (asymptotic) advantage to implementing a centralized protocol, while OED still gives an exponential improvement over greedy protocols. We believe this insight may be valuable in VLSI design where protocols like OED could be implemented at a hardware level in order to reduce buffer space requirements. Hardware implementations of similar protocols have been proposed, for example in [3], in order to achieve decentralized (gradient) clock synchronization.

2 Model and Preliminaries

We model a packet forwarding network as a directed graph, $G = (V, E)$. Each edge $e = (u, v) \in E$ has an associated buffer that stores packets in node u as they wait to cross the edge (u, v) to v . We use the notation $e = (u, v)$ to denote both the edge in G and its associated buffer.

In our model, an execution proceeds in synchronous rounds. Each round consists of two steps: an **injection step** in which new packets arrive in the network, and a **forwarding step** in which buffers forward packets across edges of the graph. During the forwarding step, each buffer chooses a subset of packets to forward, and forwards those packets across the edge (u, v) associated with the buffer. These packets arrive at their next location – either another buffer in node u , or are delivered to their destination – before the beginning of the next round. Each edge e , has a **capacity** $C(e)$, which is the maximum number of packets that can cross e in a single forwarding step.

At a given time t , we use $L^t(e)$ to denote the contents of buffer e during round t between the injection and forwarding steps. $|L^t(e)|$ is the **load** of e – i.e., number of packets stored in buffer e .

A **packet** p is a pair (t, P) where $t \in \mathbf{N}$ and $P = (v_0, v_1, \dots, v_\ell)$ is a directed path in G . The interpretation is that t indicates the time (round) at which P is injected, and P specifies a **route** from P 's **source**, $e_0 = (v_0, v_1)$ to P 's **destination** v_ℓ . An **adversary** or **injection pattern** A is a multi-set of packets.

Given a packet $p = (t, (v_0, v_1, \dots, v_\ell))$, we say that p 's route **contains** an edge $e \in E$ if $e = (v_i, v_{i+1})$ for some $i \in [\ell - 1]$. For a fixed adversary A and time interval $T = [r, s] \subseteq \mathbf{N}$, we define $N^T(e)$ to be the number of packets injected during times $t \in T$ whose routes contain e . That is

$$N^T(e) = |\{(t, P) \in A \mid t \in T \text{ and } P \text{ contains } e\}|.$$

We also define a more refined measure of utilization of an edge e that differentiates packets according to their origins. Specifically, for any subset $S \subseteq E$, we define

$$N_S^T(e) = |\{(t, P) \in A \mid t \in T, (v_0, v_1) \in S, \text{ and } P \text{ contains } e\}|$$

In particular, we have $N^T(e) = N_E^T(e)$. In the adversarial queueing model (AQT) of Borodin et al. [2], the edge utilization of an adversary A is parameterized as follows.

► **Definition 2.1.** Given $\rho > 0$ and $\sigma \geq 0$, we say that an adversary A is (ρ, σ) -**bounded** if for all e and (finite) intervals $T \subseteq \mathbf{N}$ we have

$$N^T(e) \leq \rho |T| + \sigma. \quad (1)$$

We denote the family of (ρ, σ) -bounded adversaries by $\mathcal{A}(\rho, \sigma)$.

For a (ρ, σ) -bounded adversary, the parameter ρ is an upper bound on the maximum average utilization of an edge in the network, while σ measures “burstiness” – the amount by which the average can be exceeded over any time interval. For example, taking T with $|T| = 1$, (1) implies that at most $\rho + \sigma$ packets are injected into any buffer in any single round.

2.1 Locally Bursty Adversaries

Here, we define a more refined parameterization of adversaries, which we call the *local burst model*. We refer to adversaries parameterized by the local burst model as *locally bursty adversaries*, or *LBA*s.

► **Definition 2.2.** Let A be an adversary, $\rho > 0$, $\sigma \geq 0$ and $\beta : E \rightarrow \mathbf{N}$. Then we say that A is **locally** (ρ, σ, β) -**bounded** if for all finite intervals $T \subseteq \mathbf{N}$, subsets $S \subseteq E$ and $e \in E$, we have

$$N_S^T(e) \leq \rho |T| + \sigma + \sum_{f \in S} \beta(f). \quad (2)$$

That is, for every subset S of buffers, the rate of injections into S that cross e only ever exceeds ρ by σ more than the sum of the $\beta(e)$ for $e \in S$. We denote the family of local (ρ, σ, β) -bounded adversaries by $\mathcal{L}(\rho, \sigma, \beta)$. In the case that $A \in \mathcal{L}(\rho, \sigma, \beta)$ and there is a constant B such that $\beta(e) \leq B$ for all buffers e , we will say that A is *local* (ρ, σ, B) -bounded.

We formalize the following observation that gives a relationship between the parameters of (ρ, σ) -bounded adversaries and local (ρ, σ, β) -bounded adversaries.

► **Observation 2.3.** Fix a network G and parameters ρ , σ , and $\beta : E \rightarrow \mathbf{N}$. Suppose $A \in \mathcal{L}(\rho, \sigma, \beta)$. Then $A \in \mathcal{A}(\rho, \sigma')$ for $\sigma' = \sigma + \sum_{e \in E} \beta(e)$.

► **Example 2.4.** Let G be the *single-destination path* of size n . That is, $G = (V, E)$ where $V = \{1, 2, \dots, n, n+1\}$ and $E = \{(i, i+1) \mid i \in [n]\}$. Further, all injected packets have destination $n+1$. We consider two injection patterns, A_0 and A_1

A_0 : in rounds $1, n+1, 2n+1, \dots, kn+1, \dots$, there are n packets injected into buffer 1 with destination $n+1$.

A_1 : in rounds $1, n+1, 2n+1, \dots, kn+1, \dots$, one packet is injected into each buffer $i = 1, 2, \dots, n$ with destination $n+1$.

Observe that both adversaries are in $\mathcal{A}(1, n-1)$, but not in $\mathcal{A}(1, \sigma)$ for any $\sigma < n-1$. Thus, the parameters of Definition 2.1 do not distinguish A_0 and A_1 . Yet A_0 and A_1 have vastly different buffer space requirements. A_0 requires buffer 1 to have space n for any forwarding protocol, while simple greedy forwarding for A_1 will achieve buffer space usage $|L^t(i)| \leq 1$ for all t and i .

The parameters of the local burst model, however, can distinguish between A_0 and A_1 . $A_1 \in \mathcal{L}(1, 0, 1)$ (i.e., $\beta(i) = 1$ for all i), while $A_0 \in \mathcal{L}(1, \sigma, \beta)$ only for $\sigma + \beta(1) \geq n-1$. We will show that in the case of *information gathering networks* – networks in which all

packets share a common destination and the union of their routes forms a tree – all local $\mathcal{L}(1, \sigma, B)$ -bounded adversaries can be forwarded using $O(B \log n + \sigma)$ space. Thus, the *local* burst parameter β gives a more refined understanding of the buffer space requirement of a given injection pattern.

2.2 Flows

Another well-studied model for packet forwarding is the network calculus model introduced by Cruz [4, 5]. In the network calculus, packets are associated with *flows*, and their arrivals are modeled as continuous time processes.

► **Definition 2.5.** *Given a network $G = (V, E)$, A **flow** $\phi = (a, P)$ consists of a right-continuous arrival curve $a : \mathbf{R} \rightarrow \mathbf{R}$ and associated path P . We say that ϕ has **rate** (at most) r and burst parameter b if for all $s < t$, the arrival curve a satisfies*

$$a(t) - a(s) \leq r \cdot (t - s) + b. \quad (3)$$

By convention, we assume $a(t) = 0$ for all $t < 0$.

For a single flow ϕ , the parameters r and b are analogous to the rate and burst parameters ρ and σ in Definition 2.1. However, in a flow ϕ , all packets share a common route, P . In particular, all packets associated with ϕ are injected to the same buffer and have the same destination.

In order to consider scenarios in which packets have multiple routes, we must consider multiple concurrent flows. In this setting, the analogy between equations (1) and (3) breaks down, as the former bounds the total number packets utilizing any particular edge, while the latter bounds the arrivals of packets in flows (i.e., along entire paths, rather than individual edges). In order to tighten the connection between the AQT injection model and flows, we introduce a **dependent** flow model in which we constrain the sum of arrival rates of flows across edge.

► **Definition 2.6.** *Let $G = (V, E)$ be a network and $\Phi = \{\phi\}$ be a family of flows. For an edge $e \in E$, let Φ_e denote the set of flows in Φ whose paths contain e . That is,*

$$\Phi_e = \{(a, P) \in \Phi \mid e \in P\}.$$

*Suppose each $\phi \in \Phi$ obeys a r_ϕ, b_ϕ bound as in (3). We say that Φ obeys a **locally dependent rate bound** r and **global burst parameter** σ if for every edge e , every set $\Psi \subseteq \Phi_e$ of flows, and all times s, t , we have*

$$\sum_{\phi \in \Psi} (a_\phi(t) - a_\phi(s)) \leq r \cdot (t - s) + \sigma + \sum_{\phi \in \Psi} b_\phi. \quad (4)$$

We note the similarity between equations (4) and (2). In fact, Definition 2.6 is a strict generalization of the LBA model: Given any injection pattern A , we can associate a family Φ_A of flows with A . Specifically, we define Φ_A to be

$$\Phi_A = \left\{ (a, P) \mid (t, P) \in A \text{ and } a(t) = \sum_{s \in \mathbf{N}, s \leq t} |\{(s, P) \in A\}| \right\} \quad (5)$$

With this association, the following lemma is clear.

► **Lemma 2.7.** *Suppose A is a locally (ρ, σ, β) -bounded adversary, and let Φ_A be the corresponding flow defined by (5). Then for each flow $\phi = (a_\phi, P_\phi) \in \Phi_A$, a has rate at most ρ , global burst parameter σ , and local burst parameter $b_\phi = \beta(\text{ini}_\phi)$, where ini_ϕ denotes the initial buffer in ϕ 's path. Moreover, Φ_A obeys a locally dependent rate bound of ρ .*

Conversely, LBAs arise naturally as discretizations (packetizations) of (locally dependent) flows. We formalize this connection in the following definition and lemma.

► **Definition 2.8.** *Let $G = (V, E)$ be a network and Φ a family of flows on G . The **discretization** of Φ is the AQT injection pattern A_Φ defined as follows. For each flow $\phi = (a_\phi, P_\phi) \in \Phi$ and time $t \in \mathbf{N}$, A_Φ contains $\lfloor a_\phi(t) \rfloor - \lfloor a_\phi(t-1) \rfloor$ packets injected at time t with route P_ϕ .*

We can view the discretization of a flow as forming packets via the following process. Each buffer maintains a set of (complete) packets, as well as a reserve of “fractional” packets associated with each flow originating at the buffer. At times $s \in (t-1, t]$, flows enter a buffer e . At time t , the integral parts of each flow that has not yet been bundled as packets are injected as complete packets into the buffer, while the fractional remainder is reserved. The following lemma shows that for flows obeying a locally dependent rate bound, the resulting packet injection pattern is locally bounded as well.

► **Lemma 2.9.** *Let $G = (V, E)$ be a graph and Φ a family of flows on G . For each $\phi \in \Phi$, let ini_ϕ denote the initial buffer in ϕ 's path. Suppose Φ obeys a locally dependent rate bound of r with global burst parameter σ , and define the function $\beta : E \rightarrow \mathbf{N}$ by*

$$\beta(e) = \sum_{\phi : \text{ini}_\phi = e} (1 + b_\phi).$$

Then the discretization A_Φ of Φ is locally (r, σ, β) bounded.

Proof. Fix a set $S \subseteq E$ of initial buffers, an edge e , and (discrete) time interval $T = [t_0, t_1]$. Let $A = A_\Phi$, and let $\Psi \subseteq \Phi_e$ be the subset of flows containing e and with origin in S . We compute

$$\begin{aligned} N_S^T(e) &= |\{(t, P) \in A \mid t \in T, (v_0, v_1) \in S, \text{ and } P \text{ contains } e\}| \\ &= \sum_{\phi \in \Psi} \sum_{s \in T} (\lfloor a_\phi(s) \rfloor - \lfloor a_\phi(s-1) \rfloor) \\ &= \sum_{\phi \in \Psi} (\lfloor a_\phi(t_1) \rfloor - \lfloor a_\phi(t_0-1) \rfloor) \\ &\leq r \cdot (t_1 - t_0 + 1) + \sigma + \sum_{\psi \in \Phi} (1 + b_\psi) \\ &= r \cdot |T| + \sigma + \sum_{e \in S} \beta(e). \end{aligned} \tag{6}$$

In Equation (6), we use the fact that $\lfloor a \rfloor - \lfloor b \rfloor \leq 1 + a - b$. ◀

► **Remark 2.10.** The result of Lemma 2.9 is a significant refinement of the analogous statement for the standard (ρ, σ) burst model. To see this, consider the single destination path (Example 2.4), and take $\Phi = \{\phi_1, \phi_2, \dots, \phi_n\}$ to be the family of flows where each ϕ_i has arrival curve $a(t) = \frac{1}{n}t$ and associated path $P_i = (i, i+1, \dots, n+1)$. In the associated injection pattern A_Φ , one packet is injected into every buffer at times $n, 2n, 3n, \dots$ (cf. A_1 in Example 2.4). Even though flows in Φ have burst parameter 0, large bursts appear in A_Φ as the result of the rounding process. Nonetheless, Lemma 2.9 asserts that A_Φ is *locally* $(1, 0, 1)$ -bounded, while the injection pattern is only $(1, \sigma)$ -bounded for $\sigma \geq n-1$.

3 Packet Bundling

In this section we assume that for a network $G = (V, E)$, all edges have the same (integral) capacity C . We examine the following strategy for dealing with general uniform capacity networks: when packets arrive in a buffer, they are set in a reserve buffer until sufficiently many (e.g., C) packets occupy the reserve buffer. Then the packets are bundled together, and treated as one indivisible “jumbo” packet. This process is appealing because if all jumbo packets have size C , then forwarding protocols designed for unit capacities can be applied to jumbo packets. Thus, the approach sidesteps potential subtleties in reasoning about general capacities (see Section 1.1 in [6]).

Our main results in this section show that if the original packet injection pattern obeys an LBA bound, then the resulting injection pattern of jumbo packets obeys a similar LBA bound with the parameters scaled down. Thus, if any algorithm guarantees some buffer space usage for unit capacity networks, then applying the same algorithm to jumbo packets will automatically give an analogous bound for general capacities.

3.1 Uniform Packets

We first consider the case where all packets have unit size (as in the standard AQT model), but all edges in the network have capacity $C \geq 1$. Now let A be any locally (C, σ, β) -bounded adversary, and let $\Phi = \Phi_A$ be the corresponding family of flows. We define the C -*reduction* of Φ , denoted $\frac{1}{C}\Phi$, to be

$$\frac{1}{C}\Phi = \left\{ \left(\frac{1}{C}a, P \right) \mid (a, P) \in \Phi \right\}.$$

Similarly, we define the C -reduction of A , denoted $\frac{1}{C}A$, to be the discretization (Definition 2.8) of $\frac{1}{C}\Phi$.

Observe that $\frac{1}{C}A$ is derived from A via precisely the process of forming jumbo packets as described above. The following proposition follows immediately from Lemmas 2.7 and 2.9.

► **Proposition 3.1.** *Suppose A is locally (ρ, σ, β) -bounded. Then, $\frac{1}{C}A$ is locally $(\rho/C, \sigma/C, 1 + \beta/C)$ -bounded.*

Again, we emphasize that the analogue of Proposition 3.1 is not true for the standard (ρ, σ) -bounded adversary model. The proposition has the following consequence.

► **Corollary 3.2.** *Suppose F is a forwarding protocol that for any locally $(1, \sigma, \beta)$ -bounded adversary A on a unit-capacity network $G = (V, E)$ achieves buffer space usage*

$$\sup_{e \in E, t \in \mathbf{N}} |L^t(e)| \leq f_G(\sigma, \beta).$$

Then for any uniform capacity C and locally (C, σ, β) -bounded adversary A , applying F to $\frac{1}{C}A$ achieves buffer space usage

$$\sup_{e \in E, t \in \mathbf{N}} |L^t(e)| \leq C f_G(\sigma/C, 1 + \beta/C) + C.$$

We note that the additive C term in the final expression comes from the need to store packets that have not yet been bundled.

3.2 Heterogeneous Packets

The framework described in Sections 2.2 and 3.1 shows how forwarding protocols for the AQT model with unit edge capacities can be applied to (1) discretizations of continuous flows, and (2) AQT adversaries with arbitrary uniform edge capacities and (uniform) unit-sized packets. Here, we describe a slight modification of the framework that allows for indivisible packets with heterogeneous sizes. To this end, we augment the AQT model as follows:

- Each packet p has an associated size, denoted $w(p)$.
- In a single round, an edge with capacity C can forward a set of packets whose sizes sum to at most C .
- An adversary A is locally (ρ, σ, β) bounded if for any subset S of buffers, any edge e , and in any T consecutive rounds, the sum of *sizes* of packets injected into S whose paths contain e is at most $\rho \cdot T + \sigma + \sum_{f \in S} \beta(f)$.

The following example shows one complication caused by indivisible heterogeneous packets.

► **Example 3.3.** Consider a single edge e with capacity 1. Then a $(1, 1)$ -bounded adversary can inject 3 packets of size $2/3$ every 2 rounds that must cross e . Since e has capacity 1, it can only forward a single packet each round. Thus, the injection pattern is infeasible (i.e., cannot be handled with finite buffer space).

We can preclude infeasible injection patterns (such as Example 3.3) by further restricting the allowable injection rate. Consider the following bundling procedure: when packets are injected into a buffer, they are placed in a reserve buffer. If the load of the reserve buffer exceeds $\frac{1}{2}C$, then its contents are bundled into packets, each of whose total load is at least $\frac{1}{2}C$. Arguing as before, if the original injection pattern is locally (ρ, σ, β) -bounded with $\rho \leq \frac{1}{2}C$, then the resulting injection pattern of bundled packets is locally $(1, \sigma/C, 1 + \beta/C)$ -bounded. Note that even though the rate of the adversary is $\rho \leq 1/2$, the rate of the bundled injections can be as large as 1. This occurs, for example, if the adversary always injects $C/2$ packets into a single buffer each round. These packets are then bundled, resulting in one complete bundle appearing each round.

4 OED Upper Bound

In this section we prove worst-case buffer space upper bounds for “information gathering networks” with a locally bursty adversary – i.e., instances in which all packets share a common destination and the union of trajectories of all packets forms a tree. Specifically, we show that the odd-even downhill (ODE) algorithm of [6, 10] requires $O(B \log n + \sigma)$ buffer space for any (ρ, σ, β) -bounded adversary for which $\beta(e) \leq B$ for all buffers e . For clarity and notational simplicity we describe the algorithm and argument in the simpler setting where the network consists of a path. All of the results remain true for general information gathering networks, and analogous arguments follow using the terminology and preliminary results described in [10].

In the case of the single destination path, the network $G = (V, E)$ consists of a path: $V = \{1, 2, \dots, n + 1\}$, and $E = \{(i, i + 1) \mid i \leq n\}$. All packets share the destination $n + 1$, though they can be injected into any buffer. To cut down on notational clutter, we associate a buffer $(i, i + 1)$ with its index i . In this setting, we describe the **Odd-Even Downhill** or **OED** algorithm independently introduced by Dobrev et al. [6] and Patt-Shamir and Rosenbaum [10]. Following Section 3, we assume that all edge capacities are 1, and that all packets have unit size. To simplify notation, we use $L^t(i)$ to denote the the *number* of packets in buffer i immediately before the forwarding step of round t .

► **Definition 4.1.** *The OED rule stipulates that i forwards a packet in round t if and only if one of the following conditions is satisfied:*

1. $L^t(i) > L^t(i + 1)$, or
2. $L^t(i) = L^t(i + 1)$ and $L^t(i)$ is odd.

By convention, we set $L^t(n + 1) = 0$ for all t .

Both original papers [6, 10] show that for all (ρ, σ) -bounded adversaries, the maximum buffer load under OED forwarding is $O(\log n + \sigma)$. We will show that OED forwarding achieves similar buffer space usage for any local (ρ, σ, β) -bounded adversary.

► **Theorem 4.2.** *Let G be a single-destination path of size $n + 1$, and let A be any local (ρ, σ, β) -bounded adversary where $\beta(i) \leq B$ for all i . Then the worst case buffer load is $O(B \log n + \sigma)$. That is,*

$$\sup_{t,i} L^t(i) = O(B \log n + \sigma).$$

Our proof of Theorem 4.2 follows the analysis of OED presented in [10]. Specifically, their analysis considers the evolution of *plateaus* in the network.

► **Definition 4.3** (cf. [10]). *Let G be a single destination path and $L^t : V \rightarrow \mathbf{N}$ a **configuration** – i.e., assignment of loads to buffers – at some time t . We say that an interval $I = [a, b] \subseteq V$ is a **plateau** of **height** h if I is a maximal sub-interval of V such that for all $i \in I$, $L^t(i) \geq h$. That is, every buffer in I has load at least h , and there is no larger interval I' containing I with this property. We say that I is an **even plateau** if I is a plateau of height h for some even number h .*

We think of packets in G as being arranged vertically in buffers – see Figure 1. Since packets share the same destination, for the purposes of our load analysis, we can treat all packets in a buffer as indistinguishable.¹ We refer to the **height** of a packet in a buffer as one greater than the number of packets below it. Thus, for a buffer with load 1, its sole packet is at height 1; a buffer with two packets has one at height 1 and the second at height 2, etc. We say that a packet P is **above** a plateau I of height h if $\text{ht}(P) > h$. Given a configuration $L : V \rightarrow \mathbf{N}$ and a plateau I of height h , we denote the number of packets above I by $L_h(I)$. That is,

$$L_h(I) = \sum_{i \in I} (L(i) - h). \tag{7}$$


OED forwarding does not specify *which* packet is forwarded when a buffer is forwarded, and the maximum load analysis of the algorithm is independent of this choice. Nonetheless, for the purposes of bookkeeping, it will be convenient to adopt the following conventions:

1. Whenever a packet is injected or received as the result of forwarding, it occupies the highest position in its buffer;
2. When a buffer forwards a packet, the highest packet in the buffer is forwarded.

That is – for the purposes of bookkeeping – we assume that the buffers operate as LIFO (last-in, first-out) stacks.

With these conventions, we make some preliminary observations about the movement of packets in an execution of the OED algorithm.

¹ In order to analyze packet latency, one should distinguish packets by their age.



■ **Figure 1** A configuration of packets. Each column represents a buffer, labeled 1 through 10, while the vertical axis indicates heights. The load of each buffer corresponds to the height of the highest packet in the buffer; for example, $L(6) = 5$. The indicated interval $I = [4, 8]$ is a plateau of height 2. The six shaded packets sit above the plateau I , so that $L_2(I) = 6$. Corollary 4.8 states that the number of packets above I is at most $(B + 1)|I| + \sigma$ for any locally $(1, \sigma, B)$ -bounded adversary. The proof of Theorem 4.2 applies this corollary inductively to the nested sequence of even plateaus containing the buffer with maximum load to show that the maximum load is at most logarithmic in the size of the network.

► **Lemma 4.4.** *Suppose $L : V \rightarrow \mathbf{N}$ is a configuration immediately before forwarding and L' the configuration afterward. Suppose $I = [a, b] \subseteq V$ is an even plateau of height h in configuration L . Then in configuration L' , for all $i \in [a, b - 1]$, we have $L'(i) \geq h$. Thus, in L' , the interval $[a, b - 1]$ is contained in a plateau I' of height h .*

Proof. Suppose $i \in [a, b - 1]$. Then in L , we have $L(i), L(i + 1) \geq h$. Since h is even, i will not forward unless $L(i) \geq h + 1$. Therefore, $L'(i) \geq h + 1 - 1 = h$. ◀

► **Lemma 4.5.** *For any packet P , let $\text{ht}^t(P)$ denote the height of P at time t ,² and let h be an even number. If $\text{ht}^t(P) \leq h$, then for all $s \geq t$ we have $\text{ht}^s(P) \leq h$.*

Proof. By our conventions of packet movement, the height of P in a fixed buffer is unchanged until it is forwarded. Since the top packet is always forwarded, P can only be forwarded if $\text{ht}(P) = L(i)$, where i is the buffer containing P before forwarding. Since $\text{ht}(P) \leq h$ we also have $L(i) \leq h$. Since h is even, i only forwards if $L(i + 1) \leq h - 1$. Therefore, the height of P after forwarding is at most $L(i + 1) + 1 \leq h - 1 + 1 = h$. Thus, P remains at height at most h . ◀

► **Corollary 4.6.** *Suppose $L : V \rightarrow \mathbf{N}$ is a configuration and $I = [a, b]$ is an even plateau of height h . Suppose a packet P was injected at time s , and at time $t \geq s$ sits above I (i.e., P occupies a buffer $j \in I$ and $\text{ht}^t(P) > h$). Then P was injected into a buffer in I with an initial height $\text{ht}^s(P) > h$.*

² Note that this quantity is well-defined for all t (until P is delivered to its destination) by the LIFO conventions for height and packet movement.

34:12 Packet Forwarding with a Locally Bursty Adversary

Proof. Suppose that at time t , P occupies buffer $j \in [a, b]$. Since packets are only forwarded to a buffer with larger index, P was injected at some buffer $i \leq j$ at some time $t_0 \leq t$. By Lemma 4.5, for all $s \in [t_0, t]$ we have $\text{ht}^s(P) > h$, whence the second assertion of the corollary holds.

For the first assertion, we must show that $i \in I$. To this end, for each time $s \in [t_0, t]$, let $I_s = [a_s, b_s]$ denote the plateau of height h in which P is contained at time s . Thus, I_{t_0} is the plateau above which P is initially injected so that $j \in I_{t_0}$, and $I_t = I = [a, b]$. By Lemma 4.4, for all $s < t$, we have $[a_s, b_s - 1] \subseteq I_{s+1} = [a_{s+1}, b_{s+1}]$. Therefore, we have $a_s \geq a_{s+1}$ for all s . Thus, by induction, we have $a_t \leq a_{t_0} \leq j$, where the second inequality holds because $j \in I_{t_0}$. This gives the desired result. \blacktriangleleft

We now quote a lemma from [10], which bounds the number of packets above even plateaus for (ρ, σ) -bounded adversaries.

► **Lemma 4.7** (cf. Lemma 3.4 in [10]). *Let A be a (ρ, σ) -bounded adversary and suppose L is a configuration realized by OED forwarding. Suppose I is an even plateau of height h . Then*

$$L_h(I) \leq |I| + \sigma.$$

► **Corollary 4.8.** *Let A be a local (ρ, σ, β) -bounded adversary with $\beta(i) \leq B$ for all i , and suppose L is a configuration realized by OED forwarding. Suppose I is an even plateau of height h . Then*

$$L_h(I) \leq |I| + \sigma + \sum_{i \in I} \beta(i) \leq (B + 1)|I| + \sigma.$$

Proof. For any interval I , let $A_I \subseteq A$ be the injection pattern consisting of packets injected into buffers $i \in I$. Since A is locally (ρ, σ, β) -bounded, A_I is (ρ, σ, β') -bounded with

$$\beta'(i) = \begin{cases} \beta(i) & \text{if } i \in I \\ 0 & \text{otherwise} \end{cases}$$

Therefore, by Observation 2.3, A_I is (ρ, σ') -bounded, where $\sigma' = \sigma + \sum_{i \in I} \beta(i)$.

Now let I be an even plateau of height h . By Corollary 4.6, the configuration L_h (see Equation (7)) consists entirely of packets injected into I – i.e., packets injected by A_I . Since A_I is a (ρ, σ') -bounded adversary, Lemma 4.7 implies that $L_h(I) \leq |I| + \sigma'$, which gives the desired result. \blacktriangleleft

We now have all the pieces together to prove Theorem 4.2. The idea is to use Corollary 4.8 inductively to show that plateaus cannot grow too tall.

Proof of Theorem 4.2. Assume without loss of generality that B is even, and consider any configuration L attained by A . Let $i^* = \arg \max_i L(i)$ be a buffer with maximum load, and define $I_0 \supseteq I_1 \supseteq I_2 \supseteq \dots \supseteq I_\ell$ where I_j is the plateau of height j containing i^* , and $\ell = L(i^*)$.

Define m to be the maximum value such that

$$\left| I_{m(B+2)} \right| - \frac{\sigma}{B+1} \geq 1, \tag{8}$$

if such a value of $m > 0$ exists, and take $m = 0$ otherwise. Observe that for all k , we have

$$L_k(I_k) \geq \sum_{j=k+1}^{\ell} |I_j|. \tag{9}$$

Since the I_j are nested intervals, for any $k < m$ we have

$$L_{k(B+2)}(I_{k(B+2)}) \geq (B+2) \sum_{j=k+1}^m L_{j(B+2)}(I_{j(B+2)}). \quad (10)$$

Combining (10) with the result of Corollary 4.8, we find that for all $j < m$,

$$(B+2) \sum_{j=k+1}^m |I_{j(B+2)}| \leq (B+1) |I_{k(B+2)}| + \sigma. \quad (11)$$

Rearranging (11) yields that for all $k = 0, 1, \dots, m-1$ we have

$$b \sum_{j=k+1}^m |I_{j(B+2)}| - \frac{\sigma}{B+1} \leq |I_{k(B+2)}| \quad \text{where } b = \frac{B+2}{B+1}. \quad (12)$$

Note that for $k = m-1, m-2, \dots$, (12) gives

$$b |I_{m(B+2)}| - \frac{\sigma}{B+1} \leq |I_{(m-1)(B+2)}| \quad (13)$$

$$b |I_{m(B+2)}| + b |I_{(m-1)(B+2)}| - \frac{\sigma}{B+1} \leq |I_{(m-2)(B+2)}| \quad (14)$$

⋮

Combining (13) and (14), we obtain

$$(b+b^2) |I_{m(B+2)}| - (1+b) \frac{\sigma}{B+1} \leq |I_{(m-2)(B+2)}|. \quad (15)$$

Continuing in this way, a straightforward induction argument combined with the observation that $b > 1$ gives

$$b^m \left(|I_{m(B+2)}| - \frac{\sigma}{B+1} \right) \leq |I_0| = n. \quad (16)$$

By the choice of m in (8), (16) implies that

$$m \leq \log_b n. \quad (17)$$

Again, from the definition of m , taking $h = m(B+2) + 2$, we have

$$|I_h| < 1 + \frac{\sigma}{B+1}.$$

Applying Corollary 4.8, gives

$$L_h(I_h) \leq (B+1) |I_h| + \sigma \leq B + 2\sigma + 1.$$

Since there are at most $B + 2\sigma + 1$ above I_h , the load of i^* satisfies

$$\begin{aligned} L(i^*) &\leq h + B + 2\sigma + 1 \\ &\leq (B+2)m + B + 2\sigma + 3 \\ &= O(B \log n + \sigma), \end{aligned}$$

which gives the desired result. ◀

4.1 Consequences

Here, we list some consequences of the upper bound of Theorem 4.2 when applied in combination with the packet bundling procedures described in Sections 2.2 and 3. For the following results, we assume that G is an information gathering network with uniform edge capacity C .

► **Corollary 4.9.** *Suppose A is a locally (ρ, σ, β) -bounded adversary with $\rho \leq C$ and $\beta(e) \leq B$ for all $e \in E$. Then OED forwarding applied to the C -reduction of A has buffer space usage $O((B + C) \log n + \sigma)$.*

► **Corollary 4.10.** *Suppose Φ is a family of flows with locally dependent rate bound $r \leq C$, local burst parameters $b_\phi \leq B$, and global burst parameter σ . Then OED forwarding applied to the discretization of the C -reduction of Φ requires buffer space $O((B + C) \log n + \sigma)$.*

► **Corollary 4.11.** *Suppose A is a locally (ρ, σ, β) -bounded adversary with indivisible heterogeneous packet injections, and $\rho \leq C/2$. Then OED forwarding applied to the bundled injection pattern described in Section 3.2 achieves buffer space usage $O((B + C) \log n + \sigma)$.*

5 Lower Bounds on Buffer Size

In this section, we show that buffer space usage of the OED algorithm is asymptotically optimal among deterministic forwarding protocols. In Appendix A, we generalize the lower bound to randomized forwarding protocols.

► **Theorem 5.1.** *Let F be any deterministic online forwarding protocol, and let $G = (V, E)$ be a single-destination path of length n . For any B let $\beta(e) = B$ for all $e \in E$. Then there exists a local $(1, \sigma, \beta)$ -bounded adversary A such that*

$$\sup_{t,i} L^t(i) = \Omega(B \log_B n + \sigma). \quad (18)$$

► **Remark 5.2.** The lower bound of (18) is in contrast to the centralized and semi-local upper bounds of [8, 11, 9], which show that for (ρ, σ) -bounded adversaries, maximum buffer space $O(\rho + \sigma)$ (with no n dependence) is achievable. In particular, [11] gives a smooth tradeoff between the locality of a forwarding protocol and the optimal buffer space usage. Their work shows that if each node acts based on the state of its d -distance neighborhood, then $\Theta(\frac{1}{d} \log n + \sigma)$ buffer space is necessary and sufficient.³ Thus, for (ρ, σ) -bounded adversaries, the worst-case buffer space usage for a protocol generally depends on the protocol's locality (d). Together, Theorems 4.2 and 5.1 show that this is not the case for local (ρ, σ, β) -bounded adversaries, as OED – a local ($d = O(1)$) protocol – is asymptotically optimal, even compared to centralized protocols. Thus, unlike for (ρ, σ) -bounded injection patterns, non-local information does not asymptotically improve the performance against local (ρ, σ, β) -bounded adversaries in information gathering networks.

► **Remark 5.3.** In the proof of Theorem 5.1, we describe an injection pattern as defined by an *adaptive offline adversary*. That is, the choices made by the adversary are made in response to an algorithm's forwarding decisions (i.e., the current state of all buffers in the network). If the forwarding protocol is deterministic, this assumption about the adversary is without loss of generality, as the adversary can simulate the forwarding protocol and

³ In the case $d = 1$, the algorithm of [11] reduces to the OED algorithm.

construct an injection pattern in advance. In Appendix A, we will describe a (randomized) *oblivious* adversary that is unaware of the forwarding protocol being used. Nonetheless, the oblivious adversary will almost surely require buffer space usage of $\Omega(B \log n + \sigma)$ against any (centralized, randomized) online forward protocol.

Proof of Theorem 5.1. Without loss of generality, we assume that the network size is a power of $2B$, say, $n = (2B)^m$. Given a deterministic, online forwarding protocol F , we construct an adversary A_F that injects packets in phases. Specifically, A_F chooses a nested sequence of sub-intervals $I_0 \supseteq I_1 \supseteq I_2 \supseteq \dots$ and in the k^{th} phase, A_F injects packets only into I_k . Each I_k is chosen at the end of the $(k-1)^{\text{st}}$ phase depending on the loads of buffers in I_{k-1} . The k^{th} phase lasts τ_k rounds.

We set $\tau_1 = \frac{1}{2}n$ and $I_1 = [n]$. Inductively, we define $\tau_k = \frac{1}{2B}\tau_{k-1}$ and $|I_k| = \frac{1}{2B}|I_{k-1}|$. At the beginning of phase k , A_F selects I_k and injects B packets into each buffer in I_k . For $k > 1$, A_F selects I_k as follows. Let

$$I_{k-1} = I_{k-1}^1 \cup I_{k-1}^2 \cup \dots \cup I_{k-1}^{2B},$$

where the I_{k-1}^j are consecutive intervals of size $\frac{1}{2B}|I_{k-1}|$. Then A_F selects $I_k = I_{k-1}^j$ where I_{k-1}^j has the largest total load at the end of the k^{th} phase. That is, $j = \arg \max_j L(I_{k-1}^j)$. Observe that for all k we have

$$\tau_k = \frac{1}{2}|I_k|, \text{ and} \tag{19}$$

$$\tau_k = B|I_{k+1}|. \tag{20}$$

▷ **Claim.** For all k , at the end of the k^{th} phase, the total load of I_k satisfies

$$L(I_k) \geq k \left(B - \frac{1}{2} \right) |I_k|. \tag{21}$$

Proof of Claim. We argue by induction on k . For $k = 1$, at the beginning of the first phase, B packets are injected into each buffer in the network. Thus, the total load is $Bn = B|I_1|$. After $|\tau_1| = \frac{1}{2}n$ forwarding rounds, at most $\frac{1}{2}n = \frac{1}{2}|I_1|$ packets are forwarded by the last buffer in I_1 , hence the total load in I_1 is at least $(B - 1/2)|I_1|$.

For the inductive step, assume that $L(I_{k-1}) \geq (k-1)(B - 1/2)|I_{k-1}|$ at the end of the $(k-1)^{\text{st}}$ phase. By the choice of I_k , we therefore have $L(I_k) \geq (k-1)(B - 1/2)|I_k|$ at the end of the $(k-1)^{\text{st}}$ phase. At the beginning of the k^{th} phase, A_F injects $B|I_k|$ packets into the buffers in I_k , hence the load becomes at least $(kB - (k-1)/2)|I_k|$. During the $\tau_k = |I_k|/2$ rounds of the k^{th} phase, the last buffer in I_k forwards at most $\tau_k = |I_k|/2$ packets, hence the total load of I_k decreases by at most this amount. Thus we have $k(B - 1/2)|I_k|$ as desired. ◁

Applying the claim, after $k = \log_{2B} n$ phases, we have $|I_k| = 1$ and $L(I_k) = k(B - 1/2) = \Omega(B \log_B n)$. The desired result follows by adding one more round in which A_F injects σ packets into I_k .

All that remains is to show that A_F is local $(1, \sigma, \beta)$ -bounded. To this end, suppose each k^{th} phase begins in round s_k and ends in round t_k . Then injections only occur in rounds s_k . Now fix any subset S of nodes and interval $T = [s, t]$, and define j and ℓ such that $s_{j-1} < s \leq t_j$, and ℓ with $s_\ell \leq t \leq t_\ell$. For k satisfying $j \leq k \leq \ell$, define $S_k = S \cap I_k$. Then observe that in round s_k , $B|S_k|$ packets are injected into S . Therefore, the total number of packets injected into S during T is $N = \sum_{k=j}^{\ell} B|S_k| \leq \sum_{k=j}^{\ell} B|I_k|$, while the total number of rounds is $|T| \geq \max \left\{ 1, \sum_{k=j+1}^{\ell-1} \tau_k \right\}$. We bound N as follows:

$$\begin{aligned}
N &= B|S_j| + B|S_{j+2}| + \cdots + B|S_\ell| \\
&\leq B|S_j| + B|I_{j+2}| + \cdots + B|I_\ell| \\
&= B|S_j| + \tau_{j+1} + \cdots + \tau_{\ell-1} \\
&\leq B|S| + |T|.
\end{aligned}$$

The second equality comes from Equation (20). Thus, A_F is local $(1, \sigma, \beta)$ bounded. ◀

In Appendix A, we generalize the lower bound to randomized protocols.

References

- 1 Micah Adler and Adi Rosén. Tight bounds for the performance of longest-in-system on dags. In Helmut Alt and Afonso Ferreira, editors, *STACS 2002: 19th Annual Symposium on Theoretical Aspects of Computer Science, Antibes - Juan les Pins, France, March 14–16, 2002 Proceedings*, pages 88–99, Berlin, Heidelberg, 2002. Springer Berlin Heidelberg. doi:10.1007/3-540-45841-7_6.
- 2 Allan Borodin, Jon Kleinberg, Prabhakar Raghavan, Madhu Sudan, and David P. Williamson. Adversarial queuing theory. *J. ACM*, 48(1):13–38, January 2001. doi:10.1145/363647.363659.
- 3 Johannes Bund, Matthias Függer, Christoph Lenzen, Moti Medina, and Will Rosenbaum. PALS: plesiochronous and locally synchronous systems. In *26th IEEE International Symposium on Asynchronous Circuits and Systems, ASYNC 2020, Salt Lake City, UT, USA, May 17–20, 2020*, pages 36–43. IEEE, 2020. doi:10.1109/ASYNC49171.2020.00013.
- 4 R. L. Cruz. A calculus for network delay, part I: Network elements in isolation. *IEEE Transactions on Information Theory*, 37(1):114–131, January 1991. doi:10.1109/18.61109.
- 5 R.L. Cruz. A calculus for network delay. ii. network analysis. *IEEE Transactions on Information Theory*, 37(1):132–141, 1991. doi:10.1109/18.61110.
- 6 Stefan Dobrev, Manuel Lafond, Lata Narayanan, and Jaroslav Opatrný. Optimal local buffer management for information gathering with adversarial traffic. In Christian Scheideler and Mohammad Taghi Hajiaghayi, editors, *Proceedings of the 29th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA 2017, Washington DC, USA, July 24–26, 2017*, pages 265–274. ACM, 2017. doi:10.1145/3087556.3087577.
- 7 S. Kundu and S. Chattopadhyay. *Network-on-Chip: The Next Generation of System-on-Chip Integration*. CRC Press, 2018.
- 8 Avery Miller and Boaz Patt-Shamir. Buffer size for routing limited-rate adversarial traffic. In *DISC 2016: Proceedings of the 30th International Symposium on Distributed Computing, Paris, France, September 27–29, 2016*, pages 328–341. Springer, 2016. doi:10.1007/978-3-662-53426-7_24.
- 9 Avery Miller, Boaz Patt-Shamir, and Will Rosenbaum. With great speed come small buffers: Space-bandwidth tradeoffs for routing. In Peter Robinson and Faith Ellen, editors, *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, PODC 2019, Toronto, ON, Canada, July 29 - August 2, 2019*, pages 117–126. ACM, 2019. doi:10.1145/3293611.3331614.
- 10 Boaz Patt-Shamir and Will Rosenbaum. The space requirement of local forwarding on acyclic networks. In Elad Michael Schiller and Alexander A. Schwarzmann, editors, *PODC 2017: Proceedings of the ACM Symposium on Principles of Distributed Computing, Washington, DC, USA, July 25–27, 2017*, pages 13–22. ACM, 2017. doi:10.1145/3087801.3087803.
- 11 Boaz Patt-Shamir and Will Rosenbaum. Space-optimal packet routing on trees. In *2019 IEEE Conference on Computer Communications, INFOCOM 2019, Paris, France, April 29 - May 2, 2019*, pages 1036–1044. IEEE, 2019. doi:10.1109/INFOCOM.2019.8737596.

- 12 David Salyers, Yingxin Jiang, Aaron Striegel, and Christian Poellabauer. Jumbogen: Dynamic jumbo frame generation for network performance scalability. *SIGCOMM Comput. Commun. Rev.*, 37(5):53–64, October 2007. doi:10.1145/1290168.1290174.
- 13 Stefan Schmid and Jukka Suomela. Exploiting locality in distributed sdn control. In *Proceedings of the Second ACM SIGCOMM Workshop on Hot Topics in Software Defined Networking, HotSDN '13*, pages 121–126, New York, NY, USA, 2013. Association for Computing Machinery. doi:10.1145/2491185.2491198.

A Generalization to Randomized Protocols

Here, we generalize the lower bound of Theorem 5.1 to *randomized* forwarding protocols. Specifically, we construct a randomized *oblivious* injection pattern that requires buffer space $\Omega(B \log n + \sigma)$ against any online forwarding protocol.

► **Theorem A.1.** *Let $G = (V, E)$ be the single destination path of length n . Then for every σ , and B , there exists a randomized injection pattern $A_{\text{rand}} \in \mathcal{L}(1, \sigma, B)$ such that for any (centralized, randomized) online forwarding protocol F , we have*

$$\sup_{t,i} L^t(i) = \Omega(B \log_B n + \sigma)$$

almost surely.

Again, we emphasize the order of quantifiers in the statement of the Theorem A.1: the *same* randomized injection pattern achieves the lower bound (almost surely) for every forwarding protocol.

The adversary A_{rand} of Theorem A.1 is a straightforward modification of the adversary A_F constructed the proof of Theorem 5.1. Recall that A_F injects packets into a nested sequence of intervals $I_0 \supseteq I_1 \supseteq \dots \supseteq I_k$ where $k = O(\log_B n)$, and $|I_k| = 1$. Each for $j \geq 1$, I_j is chosen to be one of $2B$ sub-intervals of I_{j-1} with maximum average load and $|I_{j-1}| = 2B |I_j|$. The idea of A_{rand} is to perform the same injection pattern as A , except that A_{rand} chooses each $I_j \subseteq I_{j-1}$ randomly, independent of the choices of the forwarding protocol. We will show that for any execution of any forwarding protocol, injecting in this way yields a load of $\Omega(B \log n + \sigma)$ with probability $\Omega(1/n)$. Thus, by independently repeating the randomized injection pattern ad infinitum, the lower bound is achieved almost surely (and with high probability after $O(n^2)$ injection rounds).

In order to formalize our description of A_{rand} , we first observe that the sequence $I_0 \supseteq I_1 \supseteq \dots \supseteq I_k$ of intervals chosen by A_F is uniquely determined by $I_k = [a_k]$, the final buffer into which A_F injects packets. We also note that the injection pattern A_F consists of $O(n)$ injection rounds, in which $O(Bn + \sigma)$ packets in total are injected. Let A_i denote such an injection pattern in which $I_k = [i]$ – i.e., i is the final buffer into which A_i injects packets.

► **Definition A.2.** *Let A_i be the injection pattern described in the preceding paragraph. Then the adversary A_{rand} injects packets as follows. Repeat:*

1. choose $i \in [n]$ uniformly at random.
2. in $O(n)$ rounds, inject packets as in A_i
3. wait $Bn + \sigma$ rounds without injecting any packets

*A single iteration of step 1–3 is an **epoch**, and each epoch consists of $k = O(\log n)$ **phases** (corresponding to the sub-intervals I_j chosen in A_i).*

34:18 Packet Forwarding with a Locally Bursty Adversary

► **Definition A.3.** Consider an execution of some forwarding protocol F with adversary A_{rand} . We say that phase $j > 0$ of some epoch of A_{rand} is **good** if at the beginning of the phase we have

$$\frac{L(I_j)}{|I_j|} \geq \frac{L(I_{j-1})}{|I_{j-1}|}.$$

We say that an epoch is **good** if all of its phases are good.

The following corollary follows immediately from the proof of Theorem 5.1.

► **Corollary A.4.** Suppose an execution of a protocol F with adversary A_{rand} experiences a good epoch with injection pattern A_i . Then in the epoch's final injection round, $L(i) = \Omega(B \log n + \sigma)$.

By Corollary A.4, all that remains to prove Theorem A.1 is to show that each epoch is good with sufficient probability.

► **Lemma A.5.** Consider a single epoch of an execution of a protocol F with adversary A_{rand} . Then each phase is good independently with probability at least $1/2B$.

Proof. Consider the interval I_{j-1} at the beginning of the j^{th} phase. There are $2B$ choices of the sub-interval I_j , and each is chosen with equal probability. By the pigeonhole principle, at least one choice I_j is good. Since A_i is chosen uniformly at random with $i \in [n]$, conditioned on I_{j-1} , the choice of I_j is uniformly at random. Thus, the probability I_j is a good choice is at least $1/2B$. Finally, we note that since i is chosen uniformly at random, the choice of which sub-interval of I_{j-1} is chosen in phase j is independent of the sub-intervals chosen in other phases. ◀

► **Corollary A.6.** Each epoch is good independently with probability at least $1/n$.

Proof. Each epoch consists of $\log_{2B} n$ phases, and each phase is good independently with probability $1/2B$. Therefore, the probability that all phases are good is $(1/2B)^{\log_{2B} n} = 1/n$. By construction, the choices of i used for each epoch are mutually independent. ◀

Proof of Theorem A.1. The proof of Theorem 5.1 implies that each A_i chosen by A_{rand} is locally (ρ, σ, β) -bounded. Waiting $Bn + \sigma$ rounds with no injections between epochs ensures that A_{rand} is locally (ρ, σ, β) -bounded as well.

By Corollary A.4, if A_{rand} experiences a good epoch, then it inflicts a buffer load of $\Omega(B \log n + \sigma)$. By Corollary A.6, the probability that the first k epochs are not good is at most $(1 - 1/n)^k$. Taking the limit as $k \rightarrow \infty$, we find that $\Pr(\sup_{t, i'} L^t(i') = o(B \log n + \sigma)) = 0$. ◀

► **Remark A.7.** The proof of Theorem A.1 shows that the probability that none of the first k epochs are good is at most $(1 - 1/n)^k$. This expression implies that a good epoch occurs with high probability after $k = O(n \log n)$ epochs. Since each epoch lasts $O(Bn + \sigma)$ rounds, a good epoch (hence a load of $\Omega(B \log n + \sigma)$) occurs after $O(Bn^2 \log n)$ rounds with high probability.