

Broadcasting on Adversarial Multiple Access Channels ^{*}

Bader A. Aldawsari [†] Bogdan S. Chlebus [‡] Dariusz R. Kowalski [‡]

Abstract

We study broadcasting on multiple-access channels under adversarial packet injection. Leaky-bucket adversaries model packet injection. There is a fixed set of stations attached to a channel. Additional constraints on the model include bounds on the number of stations activated at a round, individual injection rates, and randomness in generating and injecting packets. Broadcast algorithms that we concentrate on are deterministic and distributed. We demonstrate that some broadcast algorithms designed for ad-hoc channels have bounded latency for wider ranges of injection rates when executed on channels with a fixed number of stations against adversaries that can activate at most one station per round. Individual injection rates are shown to impact latency, as compared to the model of general leaky bucket adversaries. Outcomes of experiments are given that compare the performance of broadcast algorithms against randomized adversaries. The experiments include randomized backoff algorithms.

Key words: multiple-access channel, distributed broadcast, adversarial packet injection, queue stability, packet latency.

^{*}An earlier version of this work was presented at the 18th IEEE International Symposium on Network Computing and Applications (NCA 2019) [4].

[†]Department of Computer Science and Engineering, University of Colorado Denver, Denver, Colorado, USA.

[‡]School of Computer and Cyber Sciences, Augusta University, Augusta, Georgia, USA.

1 Introduction

We consider distributed broadcast algorithms on multiple-access channels. The goal is to compare their performance with respect to latency and queue sizes when packet generation is constrained by adversarial models. We investigate how the properties of channels and kinds of adversaries and classes of algorithms interplay among themselves to impact the performance of broadcasting.

Adversarial queuing is a methodology that can capture stability of communication without any statistical assumptions about traffic. It can provide a framework for worst-case bounds on performance of deterministic distributed algorithms. This approach was proposed by Borodin et al. [17] in their study of routing algorithms in store-and-forward networks, and continued by Andrews et al. [9]. Berger et al. [15] studied engineering aspects of adversarial queuing modeling packet injection. Latency of routing in adversarial queueing was studied by Aiello et al. [2], Andrews et al. [10], and Rosen and Tsirkin [29].

Multiple-access channels model contention occurring in local-area networks utilizing Ethernet protocols, see Metcalfe and Boggs [27]. Broadcasting algorithms for multiple-access channels need to resolve contention for access and typically use randomness in their design; see Chang et al. [18], Zhou et al. [30], and a survey by Chlebus [19]. Throughput of channels using Ethernet protocols was investigated by Bianchi [16] when the network is saturated, in the sense that each node has always a packet to transmit. Kwak et al. [26] studied saturated throughput of variants of backoff broadcast on multiple-access channels.

Stability of randomized communication algorithms on multiple-access channels can be considered in the queue-free model, in which a packet gets associated with a new station at the time of injection; see Chlebus [19] for a survey of this topic. The model of a fixed set of stations each with its own queue appears to have a stabilizing effect on randomized broadcasting. This was demonstrated by Håstad et al. [24], Al-Ammal et al. [3] and Goldberg et al. [23], who investigated the range of injection rates for which the binary exponential backoff is stable, as a function of the number of stations. The efficiency of backoff-related algorithms for the batched-arrival and slack-constrained models of packet creation was studied by Bender et al. [14], Bender et al. [12], Bender et al. [13], Agrawal et al. [1], Anderton et al. [8].

Broadcasting in multiple-access channels with queue-free stations in the framework of adversarial queuing was first studied by Bender et al. [11]. Deterministic distributed broadcast performed by stations with queues was introduced in the adversarial setting by Chlebus et al. [22]. The maximum throughput in such a setting was studied in [21]; that paper demonstrated that the ultimate throughput of 1 was achievable. Anantharamu et al. [7] extended work on such throughput 1 to adversarial models with individual injection rates. Anantharamu and Chlebus [5] developed deterministic distributed broadcast algorithms for ad-hoc 1-activating channels; that work also demonstrated that no broadcast algorithm can be universal in such channels. Anantharamu et al. [6] investigated latency of adversarial broadcasting by deterministic algorithms. Hradovich et al. [25] considered deterministic broadcasting on adversarial multiple-access channels when the adversary has to continually maintain the maximum allowed injection rate. Chlebus et al. [20] studied adversarial routing in multiple-access channels subject to energy constraints.

2 Technical Preliminaries

A multiple-access channel is a broadcast network that allows for each user to transmit a message directly to all the other users. A multiple-access channel consists of a shared communication medium and communication agents using the shared medium; the agents are called *nodes* or *stations*. We consider a synchronous variant of such networks, in which an execution of a communication algorithm is structured as a series of consecutive rounds.

Shared channels can be categorized depending on whether the set of nodes is set permanently or it may change in time. A *perpetual channel* has a permanent set of some n nodes, which have fixed names to identify them; a node's *name* is a unique integer in the range $[0, n - 1]$. In an *ad-hoc channel*, there is a potentially unbounded supply of nodes, which do not have any individual identifiers. Nodes may join and leave an ad-hoc channel, while only finitely many nodes participate actively in an execution of a communication algorithm at any round.

Packets obtained by the stations are encapsulated in messages, which then are broadcast on the channel. A transmitted message contains at most one packet. The duration of a round and the size of a message are scaled such that it takes a whole round to transmit a message. This means that transmissions of messages by different stations overlap if and only if they occur in the same round.

Sharing a channel. Every station obtains a feedback from the channel at every round. The feedback is the same for each station, and does not depend on whether the station transmits a message at this round or not. In particular, a transmitting node and a pausing node receive the same feedback.

A node *hears* a message when it receives it successfully. It is the critical property of multiple-access channels how the multiplicity of concurrent transmissions affects hearing of the transmitted messages. There are the following three relevant cases of the multiplicities of transmissions at a round: no nodes transmit, exactly one node transmits, and multiple (more than one) nodes transmit. When multiple nodes transmit at a round then this is referred to as a *collision* occurring in this round.

If no node transmits at a round then the channel is *silent*, which is reflected by the corresponding *silence* feedback received from the channel by every node. If only one node transmits, then all the nodes can hear the transmitted message as the feedback they receive from the channel, including the transmitting node. If multiple nodes transmit in the same round, which creates a collision, then the effect is such that the messages interfere with one another and none can be heard by any attached node. If a collision occurs at a round then every node obtains a *collision signal* as the feedback from the channel. A round in which no message is heard on the channel is called *void*; a void round is either silent or created by a collision. Multiple-access channels come in two variants, which are determined by the feedback coming from a channel when a message is not heard. A channel *without collision detection* has the property that both silence and collision signals are identical. In a channel *with collision detection* these respective signals are different, so that the nodes can distinguish between the two cases of silence and collision.

Dynamic broadcasting. Stations participate in executing a distributed broadcast algorithm. The goal is to perform *dynamic broadcasting*, in the sense that packets get generated and then injected into the nodes continually, while the nodes strive to have these packets heard on the channel. All the stations begin executing a given algorithm together starting from the first round.

Dynamic packet generation leads to additional categorization of shared channels. In the *queue-*

free model, each generated packet gets injected into a new station that has been previously passive. Such an injection makes the station active for the duration of handling the packet; this model was considered for example in Bender et al. [11]. In the *queuing* model, each station can handle multiple packets at the same time. The packets handled by a station get stored in a private buffer space referred to as this station’s *queue*; this model was considered for example in Håstad et al.[24]. In this paper we consider the queueing model. Every station has a potentially unbounded buffer to store queued packets. A packet queued by a station gets dequeued immediately after it has been heard on the channel. Packets do not get dropped by stations without being heard on the channel, neither due to timeouts nor for any other reason.

Let the nodes that have packets to transmit in the beginning of a round be called *active* in this round and *passive* otherwise. The event of injecting a packet into a passive station results in an *activation* of this station. If a positive integer k is an upper bound on the number of stations that can be activated at a round, then the communication channel is *k-activating*.

In this paper, we consider a perpetual channel with some number n of stations and adversaries constrained by the 1-activation constraint. The motivation for having at most one node activated at a round comes from the real-world applications, when new packets typically get injected into nodes that are already active, and rarely a passive node will get activated by obtaining packets to broadcast. This latter interpretation means also that having multiple passive nodes activated at a round occurs so rarely that it can be disregarded without distorting the performance of broadcasting, as modeled in simulations.

Historically, perpetual channels without constraints on the number of activations at a round were typically studied. Having a constant number of named stations allows to develop deterministic broadcast algorithms with bounded latency for all injection rates $\rho < 1$ and stable algorithms for injection rate $\rho = 1$, see work by Chlebus et al. [22, 21] and Anantharamu et al. [6]. Restricting packet injections to 1-activating patterns allows to develop deterministic broadcast algorithms for dynamic channels that are stable for sufficiently small injection rates, as was shown in Anantharamu et al. [5].

Categories of broadcast algorithms. A broadcast algorithm is *full sensing* when all the nodes listen to the channel at all times. We understand “listening to the channel” as undergoing state transitions determined by the feedback from the channel, which is opposed to “ignoring the feedback” by idling in an initial state. An algorithm is *activation based* when a passive station ignores the feedback from the channel and starts participating at a round when it gets activated by having a packet injected into it. An active station executing an activation-based algorithm listens to the channel when it has packets to transmit, but it stops listening to the channel as soon as it has no pending packets. An activation-based algorithm is called *acknowledgement based* when a station resets its state to initial after a successful transmission of a message.

A message transmitted on the channel includes at most one packet but it may consist of only control bits. If an algorithm does not use control bits at all, then messages transmitted in the course of its execution contain only packets; such an algorithm is called a *plain packet* one. A node executing a plain-packet algorithm cannot transmit a message at all if it does not have a pending packet in its queue. A general algorithm that allows stations to transmit messages with control bits is also called *adaptive*. A node executing an adaptive algorithm may transmit a message consisting of only control bits.

The actions performed at a round by a node executing a full-sensing algorithm or an activation-

based one by an active station are as follows. The node first either transmits a message or pauses, as determined by its state. Then the node obtains feedback from the channel, which is the same at all the stations. Next, the node may have a number of packets injected into it at this round. Finally, the node performs local computation, which can be interpreted as a state transition. Local computation involves the following actions. If packets were injected then they are enqueued in the node's queue. The private variables get updated, depending on what occurred in the round up to this moment. The node decides if to transmit at the next round, and if so then it builds a message to be transmitted. In particular, if the queue includes packets then the message to be transmitted may include a packet.

Performance metrics. The number of packets in a station's queue at a round is this station's *queue size* at the round. If a packet is injected into a station at a round t_1 and is heard on the channel at a round $t_2 > t_1$ then $t_2 - t_1$ is the number of rounds it spends in the queue, which is this packet's *delay*.

The performance of broadcast algorithms is assessed with respect to two performance metrics. *Latency* is the maximum packet delay in an execution, and the *number of queued packets* or simply *queues* is the maximum sum of the queue sizes of the stations at a round in an execution. A broadcast algorithm is *stable* in an execution if there is an upper bound on the number of queued packets.

3 Adversarial Models of Packet Injection

The quantitative restrictions on packet generation and injections are formulated as adversarial models. The adversaries we consider are all specializations of the leaky-bucket concept. The behavior of an adversary can be understood as an *execution*, which proceeds through consecutive rounds. At each round of such an execution, the adversary first generates packets and next injects these packets into stations. *Generation* means producing a finite number of packets at a round. *Injection* denotes enqueueing each generated packet into a queue at some station.

Adversarial models provide a framework that allows to study worst-case performance of broadcasting, but we can build randomness into an adversarial model to capture average performance as well. An adversary that has either the process of generation of packets or their injection determined randomly is called *randomized* and otherwise it is *worst case*. The leaky-bucket adversarial paradigm will be considered for both the worst-case and randomized-case senses. An adversary may have an upper bound on the number of stations activated at a round as an additional parameter, in particular, it could be k -activating, for an integer $k > 0$. Similarly, an adversary may be additionally constrained by upper bounds on the frequency of packets injected into each individual station.

Leaky-bucket adversaries. A leaky-bucket adversary is determined by a pair of numbers. One is the *injection rate*, denoted ρ , which is a real number satisfying $0 < \rho \leq 1$. The other is the *burstiness component*, denoted β , which is a real number satisfying $\beta \geq 1$. Together they make a *type* (ρ, β) of the adversary. For a contiguous time interval τ of $|\tau|$ rounds, the adversary may generate up to $\rho \cdot |\tau| + \beta$ packets during the rounds in τ .

The definition of the adversary of type (ρ, β) constraints both average numbers of created packets in large intervals and also bursts of numbers of packets generated in short intervals. An

injection rate ρ can be interpreted as the average number of generated packets, if averaging over large intervals. Indeed, the average number created in intervals of t rounds is at most the following: $\frac{\rho \cdot t + \beta}{t} \rightarrow \rho$, where $t \rightarrow \infty$. The adversary may generate at least one packet at a round and there is an upper bound on the number of packets created in one round. The maximum number of packets that can be generated in one round is the *burstiness* of the adversary. Indeed, let τ be a time interval of just one round: $|\tau| = 1$. Then at most $\rho \cdot |\tau| + \beta = \rho + \beta$ new packets can be generated in τ , and the following inequalities hold: $1 < \rho + \beta \leq 1 + \beta$. A generated packet is immediately injected into a station by enqueueing it in the station's queue. The adversary also determines where to inject each newly generated packet, which may be constrained though a further stipulation of the adversarial model.

The four step bucket process. Next, we describe the process of controlling the number of packets injected at a round by a leaky-bucket adversary of a given type through the concept of a bucket. We refer to the following process as a *four-step bucket process of type (ρ, β)* .

The four-step process proceeds as an execution through a sequence of rounds. The actions at a round are determined as follows. Let D be a variable, which we call a *bucket*. The bucket variable D is initialized to β at round zero. The number of packets generated in each of the following rounds is determined in four steps:

1. First, reset the bucket to $D \leftarrow \min[D + \rho, \beta]$.
2. Then, choose a non-negative integer X .
3. Next, set the number of actually generated packets to $j \leftarrow \min\{\lfloor D \rfloor, X\}$.
4. Finally, update the bucket variable to $D \leftarrow D - j$.

The first step represents the bucket *leaking* at rate ρ , which is the rate with which the bucket's available capacity grows. The second step involves a proposed quantity X of packets to be generated, assuming this quantity is consistent with the adversarial model. The third step represents a verification of the available capacity of the bucket $\lfloor D \rfloor$, as determined by the recent generations of packets and continuous leaking. The fourth step updates the available capacity of the bucket to account for the number j of packets created in this round.

Proposition 1 *Consider a sequence of packets generated at each round by an adversary in an unbounded execution. The sequence is consistent with a leaky-bucket adversarial type (ρ, β) if and only if it is consistent with the four-step bucket process of type (ρ, β) .*

Proof: Let us consider the four-step bucket process. We may observe that the invariant $0 \leq D \leq \beta$ holds after each round. We show this by induction on the round numbers. The basis of induction follows from the initialization of D to β . Assume that the invariant holds at a round and consider the next round. The bucket will satisfy $D \leq \beta$ in the next round because the first step of generation is the only one when D could be increased. The bucket will satisfy $D \geq 0$ in the next round because its decrease in the fourth step is preceded by a verification in the third step.

Now, we are ready to show that if a sequence of numbers of packets generated at each round is consistent with the four-step bucket manipulation of type (ρ, β) then it is consistent with a leaky-bucket adversary of type (ρ, β) . Let D be the bucket variable and let $\tau = [t_1, t_2]$ be a time interval.

The variable D is at most β when round t_1 begins, by the invariant. The bucket D is incremented by at most ρ in the beginning of each round in τ , for the total of at most $\rho \cdot |\tau|$ during all the rounds in τ . This can be at most balanced by injecting packets, suitably constrained by updating D in the fourth step. Generating each packet results in decrementing the variable D to $D - 1$ in the fourth step. Such a decrement is always possible to perform, because during the third step, it is verified that the number of packets generated at a round does not surpass D . It follows that the adversary can inject at most $(t_2 - t_1)\rho + \beta$ packets during τ .

Next, we show that if a sequence of numbers of packets generated at each round is consistent with a leaky-bucket adversary of type (ρ, β) then it is consistent with the four-step bucket manipulation of type (ρ, β) . Let $\tau = [0, t]$ be a time interval, for $t > 0$. The bucket variable is initialized to β at round 0, which allows to inject $\lfloor \beta \rfloor$ packets in the first round. The bucket variable is an upper bound on the number of packets that can be injected at each round, per the third step, and is decreased only to record actual packet generation in the fourth step. So the total number of packets generated in τ is accounted for by all the decrements of D . \square

Randomized adversaries. We define *randomized leaky-bucket adversaries* of type (ρ, β) , where $0 < \rho \leq 1$ and $\beta \geq 1$. The adversarial model is determined through the four-step bucket manipulation process.

Any sequence of numbers of generated packets is consistent with the definition of a leaky-bucket adversary of type (ρ, β) , by Proposition 1. The adversary is defined by how the number of packets X in the third step of packet generation is determined. We treat X as a random variable. In this randomized adversarial model. A random variable X has a Poisson distribution with parameter $\lambda > 0$ if $\Pr(X = i) = e^{-\lambda} \cdot \frac{\lambda^i}{i!}$, for each integer $i \geq 0$; see Mitzenmacher and Upfal [28].

The adversary uses a random variable X that is a Poisson distribution with parameter λ equal to the injection rate ρ in type (ρ, β) . The average number of generated packets through an execution of actions of a randomized adversary of type (ρ, β) is less than ρ . This is because the expectation of the random variable X is the second step of generation is exactly ρ , but the number of generated packets is restrained by the bucket's capacity, as reflected in the third step.

Proposition 2 *Each sequence of numbers of generated packets during a time segment $[t_1, t_2]$, for $0 \leq t_1 < t_2$, that is consistent with the general adversary of type (ρ, β) , can be generated with a positive probability by the randomized adversary of type (ρ, β) .*

Proof: Proposition 1 states that the consistence with adversary of type (ρ, β) is equivalent to the four-step bucket process of type (ρ, β) , so we can consider the corresponding bucket process. For each value D such that $0 \leq D \leq \beta$, and for each integer j such that $0 \leq j \leq D$, the probability of generating j packets in a given round is positive, as determined by the second step of the generation. \square

Individual injection rates. The adversaries with individual injection rates are defined for a perpetual channel with a given number n of stations attached to it. The stations are identified by their names in the interval $[0, n - 1]$. The adversarial model is a specialization of the general leaky-bucket adversaries, so the adversary is of some leaky-bucket adversary's type (ρ, β) .

On top of the general leaky-bucket constraint, there are additional restrictions on injecting packets into individual stations. Namely, each station i has its *individual injection rate* ρ_i assigned

to it, such that $0 \leq \rho_i \leq 1$. All these injection rates satisfy $\sum_{i=0}^{n-1} \rho_i = \rho$. For each station i and a contiguous time interval τ of $|\tau|$ rounds, the adversary can inject at most $\rho_i \cdot |\tau| + \beta$ packets into the station i during the rounds in τ . These packets are chosen from among all the packets that can be generated during τ contiguous rounds according to the general leaky-bucket constraint; there are at most $\rho \cdot |\tau| + \beta$ such packets.

4 Broadcast Algorithms

There are two general paradigms to structure deterministic broadcast in a distributed manner. Algorithms designed specifically for a perpetual multiple-access channel with a fixed set of nodes attached to the channel, each with a unique name, could operate by having the nodes exchange a token. The token visiting a node allows the node to transmit, which prevents collisions. Such algorithms could be called *token* ones; see Chlebus et al. [22, 21] and Anantharamu et al. [6]. The other paradigm is for ad-hoc channels, where nodes are dynamically generated. The idea is to assign temporary implicit names: this works for a setting in which at most one new node is added to the system at a round. Such algorithms could be called *ad-hoc* ones; see Anantharamu et al. [5].

A broadcast algorithm is *universal* if it is stable for each injection rate less than 1, and it is *strongly universal* if it is universal and there exists a range $[0, c]$, for $0 < c < 1$, such that if an injection rate ρ satisfies $0 < \rho < c$ then the number of packets in the queues is a function of only the type of the adversary (ρ, β) , rather than the number of stations n in a perpetual channel. A broadcast algorithm has *universal latency* if it has bounded latency for each injection rate less than 1. A broadcast algorithm has *strongly universal latency* if it has universal latency and there exists a range $[0, c]$, for $0 < c < 1$, such that if injection rate ρ satisfies $0 < \rho < c$ then packet delay is a function of only the type of the adversary (ρ, β) . The *throughput* is the maximum injection rate for which an algorithm is stable, if such an injection rate exists.

Broadcasting by deterministic algorithms can be more effective on perpetual multiple-access channels than on channels executing ad-hoc algorithms, due to the stabilizing effect of a fixed set of stations. In particular, a perpetual channel can achieve throughput 1, as shown in Chlebus et al. [21], and there are algorithms with universal packet latency, as shown in Anantharamu et al. [6]. In contrast to this, ad-hoc algorithms cannot handle injection rates that are at least $\frac{3}{4}$ with bounded latency, as shown in Anantharamu and Chlebus [5], which implies that there are no universal algorithms for ad-hoc channels.

We consider three groups of algorithms: deterministic token ones, deterministic ad-hoc ones, and randomized ones. What follows are outlines of specifications of the algorithms.

Algorithms for perpetual channels. Algorithm ROUND-ROBIN-WITHHOLDING (RRW) is a plain-packet full-sensing algorithm of universal latency for channels without collision detection proposed in Chlebus et al. [22]. Algorithm SEARCH-ROUND-ROBIN (SRR) is a plain-packet full-sensing algorithm of universal latency for channels with collision detection proposed in Chlebus et al. [22]. Each of the algorithms RRW and SRR can be modified by using the approach called “old-go-first,” see Anantharamu et al. [6]. This approach works as follows. An execution can be interpreted as having a token that traverses all the nodes at a round robin manner, starting from station 0. A *phase* consists of a segment of rounds in which the token makes a full cycle and returns to the starting point. The packets that are injected in a phase are considered “new” during the phase and

become “old” when the next phase starts. In the course of a phase, the new packets are ignored and only the old ones are broadcast. The algorithms modified this way are called OLD-FIRST-ROUND-ROBIN-WITHHOLDING (OF-RRW) and OLD-FIRST-SEARCH-ROUND-ROBIN (OF-SRR), respectively. Algorithm MOVE-BIG-TO-FRONT (MBTF) is an adaptive full-sensing algorithm for channels without collision detection, proposed in Chlebus et al. [21]. It attains throughput 1 and has universal latency.

Algorithms for 1-activating ad-hoc channels. Next we discuss deterministic distributed algorithms designed for 1-activating ad-hoc channels. The ad-hoc model and the algorithms were given in Anantharamu and Chlebus [5].

Algorithm COUNTING-BACKOFF is a plain-packet activation-based algorithm for ad-hoc channels with collision detection. Active nodes are stored on a virtual distributed stack, which is implemented such that an active station remembers its distance from the top of the stack.

In an execution of algorithm QUADRUPLE-ROUND, time is partitioned into segments of length four, which are considered one by one. For a processed time segment, the stations activated in one of the four rounds in the segment have an opportunity to broadcast their packets. The existence of such stations is verified in a binary search manner, with the four rounds associated with leaves. In the first round of search, each such a station is to broadcast its packet. If this first round results in silence, then this means that no station was activated in the given time segment and the algorithm advances to the next segment; otherwise, if collision is detected, then the binary search branches off to consider the left and right subtrees. If a packet is heard, then the transmitting station does not withhold the channel, instead, it will have an opportunity to transmit again in the next iteration of the binary search.

The active stations executing algorithm QUEUE-BACKOFF are organized as a distributed first-in-first-out queue, which they join in the order of activation time. A station at the front of the queue transmits its packets, and the last packet has an ‘over’ bit attached, used to prompt the next station to start its transmissions. Additionally, each message carrying a packet includes control bits describing the current size of the queue. A newly activated station transmits immediately, which results in a collision, unless the queue is empty. If a collision occurs, the front station learns that the queue has just increased, and this count is reflected in the messages the station transmits. A station that joins the queue identifies its position by the size of the queue when it learns it for the first time, from which it subtracts the number of immediately preceding collisions created by other stations joining the queue.

These three algorithms for ad-hoc channels do not use the names of nodes, even if they are available. They have ranges of injection rates with bounded latency even when there is no fixed set of nodes attached to the channel and the adversary has the power to “create” new stations by injecting packets into them, but at most one new station at a round; see Anantharamu and Chlebus [5].

Randomized backoff algorithms. Now we discuss two randomized backoff algorithms. These are the BINARY-EXPONENTIAL-BACKOFF (BEB) and QUADRATIC-BACKOFF (QB). Each of them is acknowledgement-based; as soon as a packet is made heard on the channel, the next available packet is processed immediately. A node processes its packets in the order of injection. The backoff algorithms are considered in their windowed versions. A *window* is a contiguous segment of rounds starting from the round of a successful transmission or a collision. The lengths of windows may

vary; the window just after a successful transmission is of size 1. When a new packet is available for processing, the node selects a round uniformly at random in the current window, then waits for this round to occur and transmits the packet. Since the first window consists of one round, a new packet is transmitted immediately.

Backoff algorithms differ among themselves in how the sizes of consecutive windows increase. For BINARY-EXPONENTIAL-BACKOFF, the i th window size was determined as 2^i , and for QUADRATIC-BACKOFF, the i th window size was defined to be i^2 .

The two randomized algorithms can be considered with an upper bound on the window size, as the binary exponential backoff is used in the implementation of the Ethernet. For instance, the i th window size for BINARY-EXPONENTIAL-BACKOFF could be $2^{\min(10, i)}$, which is exactly as in the Ethernet, and for QUADRATIC-BACKOFF, the i th window size could be $(\min(i, 32))^2$. Observe that, with these two choices of constants, the maximum size of a window is the same in BEB and QB, since $2^5 = 32$.

5 Upper Bounds on Queues and Latency

We investigate worst-case upper bounds on latency of adversarial broadcasting in two special cases. One is for algorithms designed for a perpetual channel with named stations when leaky bucket adversaries are restricted to have individual injection rates. The effect of individual injection rates, as compared to the general leaky-bucket adversarial model, is that the old-go-first versions of algorithms have greater bounds on packet latency rather than smaller, when compared to the regular versions. The other is for algorithms designed for ad-hoc channels when they are executed on perpetual channels.

Deterministic algorithms for 1-activating ad-hoc channels were given in Anantharamu and Chlebus [5] that can handle injection rates up to $\frac{1}{2}$ in a stable manner. A fixed set of nodes of a perpetual channel may have a stabilizing effect on these algorithms. Indeed, we show that a particular algorithm designed for ad-hoc channels has *strong* universal latency when executed on 1-activating perpetual channels.

5.1 Old-Go-First with individual injection rates

We begin with token algorithms designed for perpetual channels. It was shown in Anantharamu et al. [6] that algorithm OF-RRW executed by n stations against an adversary of type (ρ, β) has the following asymptotically-tight performance bounds: the number of packets simultaneously queued in the stations is at most $\frac{2\rho}{1-\rho} \cdot n + \beta$ and packet latency is at most $\frac{2}{1-\rho} \cdot n + \beta(1 + \rho)$. In contrast to that, algorithm RRW has the following asymptotically-tight performance bounds: the number of packets simultaneously queued in the stations is at most $\frac{2\rho}{1-\rho} \cdot n + \beta$ and packet latency is at most $\frac{2-\rho}{(1-\rho)^2} \cdot n + \frac{\beta}{1-\rho}$. It follows that algorithm OF-RRW has superior worst-case performance bounds as compared to algorithm RRW, better by a factor $1/(1 - \rho)$, which grows unbounded as ρ increases towards 1.

Next, we show that if an adversary is restricted by individual injection rates then this phenomenon does not hold.

Theorem 1 *Algorithm RRW has at most $\frac{\rho}{1-\rho}n + \beta$ packets queued and its latency is at most $\frac{2-\rho}{1-\rho}n + \beta$ when executed on a perpetual channel with n stations against a (ρ, β) adversary with individual injection rates.*

Proof: We consider an execution in which the adversary injects at full power, with the effect of burstiness considered separately. As the phases pass, their length increases approaching a limit, with possible fluctuations. If the adversary is injecting at full power into a particular station, then the number of rounds between two consecutive injections differs by at most 1, due to rounding. If this occurs for all the stations in the same phase then this phase may be extended by n rounds. Burstiness can make queues increase by β in one phase and contribute β to the length of the phase. Let p denote the number of packets in a phase in equilibrium state, while disregarding burstiness: the number of packets at the start of a phase is the same as after the phase is over. Such a phase takes $n + p$ rounds, and the number of packets injected during the phase is $\rho(n + p)$. This gives the equation $p = \rho(n + p)$, which can be solved for p to give $p = \frac{\rho}{1-\rho}n$. The number of packets in a phase can be increased by injecting up to β packets extra packets, for a total of $\frac{\rho}{1-\rho}n + \beta$.

A phase consists of the rounds spent on transmitting packets, then n silent rounds that result in the token advancing through the stations, and possibly n extra rounds due to individual fluctuations of the number of packets injected individually in each station. This bound is therefore as follows: $2n + \frac{\rho}{1-\rho}n + \beta$, which can be simplified to the claimed form by algebra. \square

Algorithm OF-SRR executed by n stations against an adversary of type (ρ, β) has the following asymptotically-tight performance bounds: the number of packets simultaneously queued in the stations is at most $\frac{4\rho}{1-\rho} \cdot n + \beta$ and packet latency is at most $\frac{4}{1-\rho} \cdot n + \beta(1 + \rho)$. In contrast to that, algorithm SRR has the following asymptotically-tight performance bounds: the number of packets simultaneously queued in the stations is at most $\frac{4\rho}{1-\rho} \cdot n + \beta$ and packet latency is at most $\frac{4-2\rho}{(1-\rho)^2} \cdot n + \frac{\beta}{1-\rho}$; see Anantharamu et al. [6].

Algorithms SRR and OF-SRR have greater upper bounds on packet latency than RRW and OF-RRW, respectively. Algorithms SRR and OF-SRR have a property that their packet latency becomes $\mathcal{O}(\log n)$ for suitably small injection rates that are $\mathcal{O}(1/\log n)$, see Anantharamu et al. [6]. Algorithm OF-SRR has superior worst-case performance bounds compared to algorithm SRR by the factor $1/(1 - \rho)$, which grows unbounded if ρ converges to 1. If an adversary is restricted by individual injection rates, then this phenomenon does not occur, as stated next.

Theorem 2 *Algorithm SRR has at most $\frac{2\rho}{1-\rho}n + \beta$ packets queued at a round and its latency is at most $\frac{3-\rho}{1-\rho}n + \beta$ if executed on a perpetual channel with n stations against a (ρ, β) adversary with individual injection rates.*

Proof: We denote by p the number of packets in a phase in equilibrium state, while disregarding burstiness: the number of packets at the start of a phase is the same as after the phase is over. There may be up to $2n - 1$ void rounds in a phase, so a phase takes at most $2n + p$ rounds, while the number of packets injected during the phase is $\rho(2n + p)$. This gives the equation $p = \rho(2n + p)$, which solved for p gives $p = \frac{2\rho}{1-\rho}n$. The number of packets in a phase can be increased by injecting up to β packets extra packets, for a total of $\frac{2\rho}{1-\rho}n + \beta$.

A phase consists of the rounds spent on transmitting packets, then $2n - 1$ silent rounds that result in the token advancing through the stations, and possibly n extra rounds due to individual

fluctuations of the number of packets injected individually in each station. This bound is therefore as follows: $3n + \frac{2\rho}{1-\rho}n + \beta$, which can be simplified to the claimed form by algebra. \square

5.2 Ad-hoc algorithms in perpetual channels

We will consider algorithms designed for ad-hoc channel when executed on a perpetual channel with a fixed set of stations against 1-activating adversaries.

Algorithm COUNTING-BACKOFF has packet latency $\frac{3\beta-3}{1-3\rho}$ in ad-hoc 1-activating channels, for injection rates ρ that are less than $\frac{1}{3}$; see Anantharamu et al. [5]. This algorithm is not stable for injection rates at least $\frac{1}{3}$, since the stack may stay forever nonempty in an execution, and the station at the bottom is starved for access to the channel. This property holds in both ad-hoc and perpetual channels.

Algorithm QUADRUPLE-ROUND is a full-sensing plain-packet algorithm for ad-hoc 1-activating channels with collision detection that has packet latency $2\beta + 4$ and queues $\beta + \mathcal{O}(1)$ against $(\frac{3}{8}, \beta)$ adversaries, see Anantharamu and Chlebus [5]. These bounds hold a fortiori for perpetual channels.

Next, we show that the algorithm attains bounded latency for a larger range of injection rates in perpetual channels compared to ad-hoc ones.

Theorem 3 *Algorithm QUADRUPLE-ROUND has queues at most $\frac{\rho}{3-7\rho} \cdot n + \beta$ and latency at most $\frac{7\rho}{(3-7\rho)^2} \cdot n + \frac{n+7\beta}{3-7\rho}$, if executed on a 1-activating perpetual channel with collision detection against (ρ, β) adversaries with $\rho < \frac{3}{7}$ and n stations attached to a channel.*

Proof: We investigate how the adversary can organize an execution to maximize the delay of some packet. The discrete time line underlying an execution is partitioned into double segments of eight consecutive rounds. It is possible to activate three stations in a segment with one packet per activated station such that algorithm QUADRUPLE-ROUND spends seven rounds processing the double segment; see Anantharamu and Chlebus [5] for similar arguments. This can be iterated with more triples of packets injected into the same stations in the double segment. Finally, one more silent round verifies that the double segment is clear. This creates the worst-case time overhead per the number of injected packets, and is the scheme repeatedly used in the argument on the level of executions.

A specific execution we present consists of two parts. In the first part, the adversary works to make all the active stations store as many packets in total as possible. This is accomplished by activating three stations in a segment, whenever there is an opportunity to do this, and such that the algorithm needs seven rounds to process such triples of packets. We keep injecting into these stations, three packets per the same stations, when the adversarial model allows to inject three packets, or a multiple of three, and when activating three new stations is not possible because fewer than three stations are passive. After the goal of the first part is accomplished, a station is activated with a new dedicated packet that we want to delay as much as possible. This is accomplished by injecting at full power only into stations that were activated prior to the dedicated packet's station.

The first part begins by the size of the set of acting stations growing with the goal to make all the available n stations active. This is feasible because injecting a packet at every other round is sufficient to maintain the current size of the queue, in the sense that a new station gets activated

immediately after a station leaves the queue. We define a *phase* to be a contiguous segment of rounds during which n stations become active. In particular, if all n stations are active simultaneously and each of them holds one packet, then a phase takes $\frac{8}{3}n$ rounds. This is because there are $\frac{n}{3}$ segments, with three stations activated in a segment, each taking 8 rounds to clear. Once the n stations are active, injecting with rate $\frac{3}{8}$ is sufficient to maintain this state in the ad-hoc channel, and with rate less than $\frac{3}{7}$ in the perpetual one.

Let p denote the maximum number of packets accrued in the first part, disregarding burstiness, so the true maximum number is $m = p + \beta$. The number of double segments in a phase is $\frac{n}{3}$, and this is the number of silent rounds resulting in clearing the double segments. When the first part of an execution is in the equilibrium state, the following equality is satisfied:

$$p = \rho \cdot \left(\frac{1}{3}n + \frac{7}{3}p \right).$$

It can be solved for p to give m : $m = \frac{\rho}{3-7\rho} \cdot n + \beta$. The delay of a packet injected now can be extended to the following value

$$\left(\frac{7}{3}m + \frac{1}{3}n \right) \left(1 + \frac{7}{3}\rho + \left(\frac{7}{3}\rho \right)^2 + \dots \right) = \frac{7m + n}{3 - 7\rho}.$$

The claimed bound is obtained from this by algebra. □

Algorithm QUEUE-BACKOFF is an adaptive activation-based algorithm for ad-hoc channels without collision detection that has packet latency $4\beta - 4$ against 1-activating adversaries with injection rate $\frac{1}{2}$, see Anantharamu and Chlebus [5].

This algorithm has strong universal latency in perpetual channels against 1-activating adversaries.

Theorem 4 *Algorithm QUEUE-BACKOFF is an adaptive algorithm with strongly-universal packet latency if executed on a 1-activating perpetual channel with collision detection. It has queues at most $\frac{\rho}{1-\rho} \cdot n + \beta$ and latency at most $\frac{\rho}{(1-\rho)^2} \cdot n + \frac{\beta}{1-\rho}$ against (ρ, β) adversaries with $\rho < 1$ and n stations attached to such a channel.*

Proof: We investigate how the adversary can organize an execution to maximize the number of packets and delay of some packet. A specific execution will consist of two parts. In the first part, the adversary will work to make all the active stations store as many packets in total as possible. This is accomplished by activating a station whenever there is an opportunity to do this and continuously injecting into this station at full power until a next station is activated. After the goal of the first part is accomplished, a station is activated with a new dedicated packet that we want to delay as much as possible. This is accomplished by injecting at full power only into stations that are in front of the dedicated packet's station in the distributed queue.

The first part begins by the queue growing to encompass all the available n stations. This is feasible because injecting a packet every other round is sufficient to maintain the current size of the queue, in the sense that immediately after a station leaves the queue at least one new station is activated. We define a phase to be a contiguous segment of rounds during which n stations become active. A phase includes n collision rounds and the rest consists of rounds with packets heard on the channel.

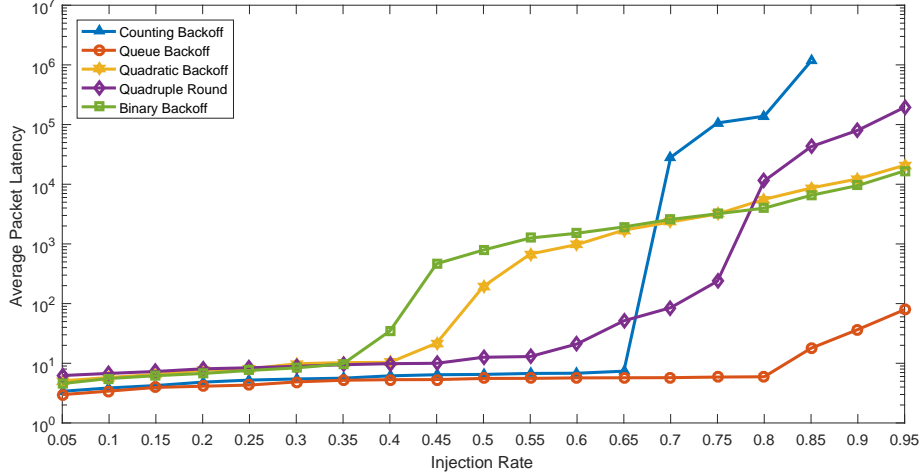


Figure 1: A comparison of ad-hoc and backoff algorithms with respect to the average packet latency for $n = 10$ stations for the full range of injection rates.

Let p be the maximum number of packets stored in all the active stations while disregarding the effect of burstiness, so that the ultimate maximum number of packets m is $m = p + \beta$. As the first part of the execution progresses, the number of packets approaches the solution of the equation $p = \rho(n + p)$. This equation can be solved for p to yield $m = \frac{\rho}{1-\rho} \cdot n + \beta$.

Once the first part ends, a packet is injected into a newly activated station with the goal to delay its successful transmission by injecting at full power in the stations preceding the packet's one. This results in the following time spent by the packet waiting:

$$m(1 + \rho + \rho^2 + \dots) = \frac{m}{1 - \rho} ,$$

which is equivalent to the claimed form by algebra. \square

Algorithm QUEUE-BACKOFF is adaptive, in that it uses control bits in messages. We conjecture that this is a necessary property of algorithms with strongly universal latency.

Conjecture 1 *There is no plain-packet algorithm that has strongly universal latency when executed on 1-activating perpetual channels.*

6 Simulations of Randomized Adversaries

We presents outcomes of experiments carried out as simulations of randomized adversaries. Such adversaries are specified by bucket processes, see Section 3. A bucket process determines the number of packets generated in a step but does not determine the stations into which the packets get injected. We specify a natural way to select stations for injecting packets that is consistent with the 1-activating constraint. Let us consider a (ρ, β) randomized adversary injecting packets into the stations of a perpetual channel with n stations. At a round, first the number of packets j to

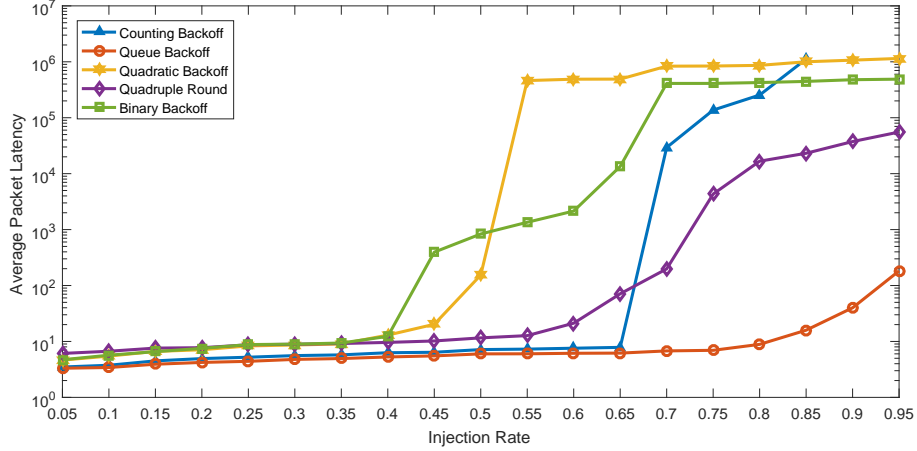


Figure 2: A comparison of ad-hoc and backoff algorithms with respect to the average packet latency for $n = 250$ stations for the full range of injection rates.

inject is determined. Next, the adversary selects one station uniformly at random from the set of passive stations as *virtually active*; if all stations are active then there is no virtually active station. The set of active and virtually active nodes makes the set of *eligible nodes* for this round. For each of the new j packets, an eligible station is assigned to it independently and uniformly at random and the packet is injected into it.

A simulating execution proceeds by running a broadcast algorithm and injecting packets subject to the type (ρ, β) of a considered adversary; in our experiments $\beta = 10$. The execution is partitioned into *stages*, such that a stage concludes with the average packet delay of a batch of some K packets; this parameter is set to $K = 5,000$ in our experiments. A stage begins by designating K consecutively generated packets as *marked* and concludes when all the marked packets have been heard on the channel. As the execution proceeds, packets are generated continuously, but if they are not marked then their delay is not measured. As a stage ends, the average packet delay is determined by the K marked packets.

We say that an average packet latency *stabilizes* when the relative average packet latencies in any two stages in four consecutive stages differ by less than 5%. If the average packet delay stabilizes then it is considered as output of the experiment, otherwise the channel is considered unstable. This assessment is conservative with respect to stability, in that some latency may be recorded even if the system is unstable in the stochastic sense. The latencies obtained in simulation may vary significantly among the algorithms, so we depict them on exponential vertical scales in the charts.

Figures 1 and 2 present a comparison of ad-hoc deterministic algorithms and acknowledgement-based randomized algorithms on the whole spectrum of injection rates and for two sample sizes of the system: one relatively small and the other larger. Latencies turn out to be comparable for small injection rates but vary significantly when injection rates approach 1. In particular, the latencies for algorithm COUNTING-BACKOFF do not stabilize in some range, which is consistent with this algorithm's instability in the worst-case sense for injection rates greater than $\frac{1}{2}$. Algorithm QUEUE-BACKOFF outperforms all the other algorithms. Individual comparisons depend on the number of

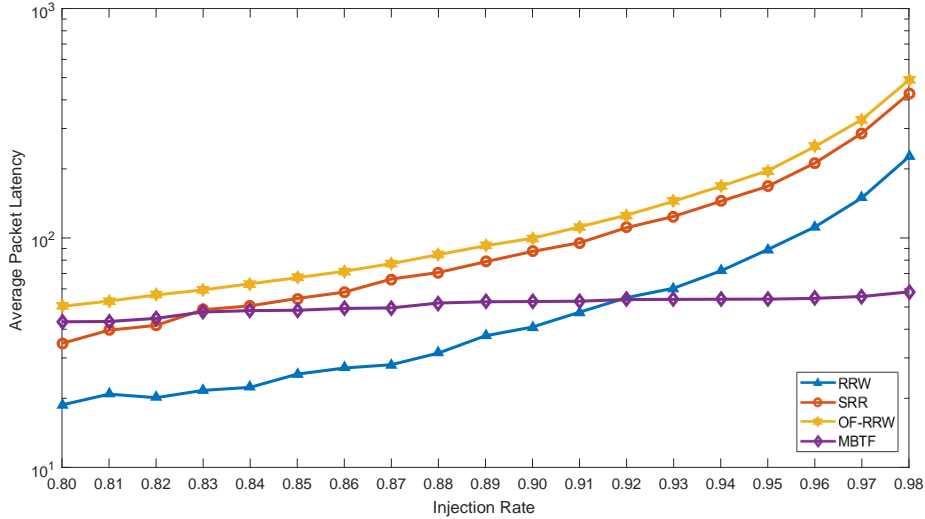


Figure 3: A comparison of average packet latency of token algorithms for the range $[0.8, 0.98]$ of injection rates with $n = 10$ stations.

stations; in particular, QUADRUPLE-ROUND may outperform both backoff algorithms, or vice versa.

Figures 3 and 4 present experiments of token algorithms for a range of large injection rates close to 1, and for two sample sizes of the system: one relatively small and the other larger. The best choice for large systems and low injection rates may be algorithm SRR. Another observation is that RRW outperforms OF-RRW in this adversarial model, which is consistent with the analysis of the individual-injection rates adversaries in Section 5.1, and the opposite of the worst-case bounds given in Anantharamu et al. [6]. Algorithm MBTF may have smallest latency for large injection rates, but when the number of stations is sufficiently small; this is consistent with the worst-case bound showed in Anantharamu et al. [6] to depend on n^2 .

7 Conclusion

We presented algorithmic contributions to the landscape of broadcasting on adversarial multiple-access channels by deterministic distributed algorithms. These results demonstrate how performance bounds of broadcast algorithms depend on model’s components. Featured examples include perpetual channels versus ad-hoc ones, and the models of unconstrained leaky-bucket injection rate versus individually-constrained injection rate.

Leaky-bucket adversaries allow to investigate stability and latency of broadcasting on multiple-access channel without stochastic assumptions on packet injection. The performance bounds of deterministic algorithms for such adversarial traffic can be worst-case only because adversarial packet injection is constrained by upper bounds on how many packets can be injected in all bounded time intervals. Simulations of such adversaries are virtually impossible because the constraints on the adversaries are in the form of an upper bound of how many packets an adversary may inject if exercising “full power of packet generation” while this behavior does not necessarily result in

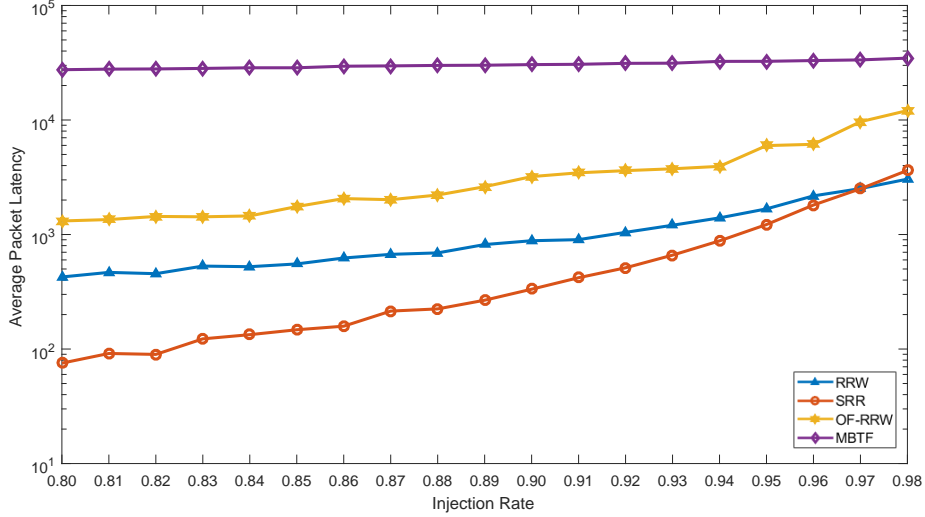


Figure 4: A comparison of average packet latency of token algorithms for the range $[0.8, 0.98]$ of injection rates with $n = 250$ stations.

worst-case performance.

We proposed a randomized model of adversarial packet injection that allows for arbitrarily bursty traffic to occur with positive probability. This model is amenable to simulations and the average performance of broadcasting algorithms can be discovered via simulated experiments. The examples of such simulations that we gave indicate a complicated landscape of behavior of broadcast algorithms, with performance bounds depending on the number of stations in a perpetual channel as well as on injection rates. These experiments involve two popular randomized backoff protocols, binary-exponential and quadratic ones. The most interesting injection rates are these close to the ultimate injection rate of 1, as they differentiate the algorithms most. Such differentiation is not necessarily simple and conclusive, as relative packet latency of two broadcast algorithms may depend on the number of stations attached to the channel. Deriving formal performance bounds on the expected packet latency and queue size of the studied broadcast algorithms for the model of randomized adversaries is an interesting direction of future research.

References

- [1] Kunal Agrawal, Michael A. Bender, Jeremy T. Fineman, Seth Gilbert, and Maxwell Young. Contention resolution with message deadlines. In *Proceedings of the 32nd ACM Symposium on Parallelism in Algorithms and Architectures (SPAA 2020)*, pages 23–35. ACM, 2020.
- [2] William Aiello, Eyal Kushilevitz, Rafail Ostrovsky, and Adi Rosén. Adaptive packet routing for bursty adversarial traffic. *Journal of Computer and System Sciences*, 60(3):482–509, 2000.
- [3] Hesham Al-Ammal, Leslie Ann Goldberg, and Philip D. MacKenzie. An improved stability bound for binary exponential backoff. *Theory of Computing Systems*, 34(3):229–244, 2001.

- [4] Bader A. Aldawsari, Bogdan S. Chlebus, and Dariusz R. Kowalski. Broadcasting on adversarial multiple access channels. In *Proceedings of the 18th IEEE International Symposium on Network Computing and Applications (NCA 2019)*, pages 1–4. IEEE, 2019.
- [5] Lakshmi Anantharamu and Bogdan S. Chlebus. Broadcasting in ad hoc multiple access channels. *Theoretical Computer Science*, 584:155–176, 2015.
- [6] Lakshmi Anantharamu, Bogdan S. Chlebus, Dariusz R. Kowalski, and Mariusz A. Rokicki. Packet latency of deterministic broadcasting in adversarial multiple access channels. *Journal of Computer and System Sciences*, 99:27–52, 2019.
- [7] Lakshmi Anantharamu, Bogdan S. Chlebus, and Mariusz A. Rokicki. Adversarial multiple access channels with individual injection rates. *Theory of Computing Systems*, 61(3):820–850, 2017.
- [8] William C. Anderton, Trisha Chakraborty, and Maxwell Young. Windowed backoff algorithms for wifi: theory and performance under batched arrivals. *Distributed Computing*, 34(5):367–393, 2021.
- [9] Matthew Andrews, Baruch Awerbuch, Antonio Fernández, Frank Thomson Leighton, Zhiyong Liu, and Jon M. Kleinberg. Universal-stability results and performance bounds for greedy contention-resolution protocols. *Journal of the ACM*, 48(1):39–69, 2001.
- [10] Matthew Andrews, Antonio Fernández, Ashish Goel, and Lisa Zhang. Source routing and scheduling in packet networks. *Journal of the ACM*, 52(4):582–601, 2005.
- [11] Michael A. Bender, Martin Farach-Colton, Simai He, Bradley C. Kuszmaul, and Charles E. Leiserson. Adversarial contention resolution for simple channels. In *Proceedings of the 17th ACM Symposium on Parallel Algorithms and Architectures (SPAA 2005)*, pages 325–332, 2005.
- [12] Michael A. Bender, Jeremy T. Fineman, Seth Gilbert, and Maxwell Young. Scaling exponential backoff: Constant throughput, polylogarithmic channel-access attempts, and robustness. *Journal of the ACM*, 66(1):6:1–6:33, 2019.
- [13] Michael A. Bender, Tsvi Kopelowitz, William Kuszmaul, and Seth Pettie. Contention resolution without collision detection. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing (STOC 2020)*, pages 105–118. ACM, 2020.
- [14] Michael A. Bender, Tsvi Kopelowitz, Seth Pettie, and Maxwell Young. Contention resolution with constant throughput and log-logstar channel accesses. *SIAM Journal on Computing*, 47(5):1735–1754, 2018.
- [15] Daniel S. Berger, Martin Karsten, and Jens B. Schmitt. On the relevance of adversarial queueing theory in practice. In *Proceedings of the International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS 2014)*, pages 343–354. ACM, 2014.
- [16] Giuseppe Bianchi. Performance analysis of the IEEE 802.11 distributed coordination function. *IEEE Journal on Selected Areas in Communications*, 18:535 – 547, 2000.
- [17] Allan Borodin, Jon M. Kleinberg, Prabhakar Raghavan, Madhu Sudan, and David P. Williamson. Adversarial queueing theory. *Journal of the ACM*, 48(1):13–38, 2001.

- [18] Yi-Jun Chang, Wenyu Jin, and Seth Pettie. Simple contention resolution via multiplicative weight updates. In *Proceedings of the 2nd Symposium on Simplicity in Algorithms (SOSA 2019)*, volume 69 of *OASICS*, pages 16:1–16:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
- [19] Bogdan S. Chlebus. Randomized communication in radio networks. In Panos M. Pardalos, Sanguthevar Rajasekaran, John H. Reif, and Jose D. P. Rolim, editors, *Handbook of Randomized Computing*, volume I, pages 401–456. Kluwer Academic Publishers, 2001.
- [20] Bogdan S. Chlebus, Elijah Hradovich, Tomasz Jurdziński, Marek Klonowski, and Dariusz R. Kowalski. Energy efficient adversarial routing in shared channels. In *Proceedings of the 31st ACM Symposium on Parallelism in Algorithms and Architectures (SPAA 2019)*, pages 191–200. ACM, 2019.
- [21] Bogdan S. Chlebus, Dariusz R. Kowalski, and Mariusz A. Rokicki. Maximum throughput of multiple access channels in adversarial environments. *Distributed Computing*, 22(2):93–116, 2009.
- [22] Bogdan S. Chlebus, Dariusz R. Kowalski, and Mariusz A. Rokicki. Adversarial queuing on the multiple access channel. *ACM Transactions on Algorithms*, 8(1):5:1–5:31, 2012.
- [23] Leslie Ann Goldberg, Mark Jerrum, Sampath Kannan, and Mike Paterson. A bound on the capacity of backoff and acknowledgment-based protocols. *SIAM Journal on Computing*, 33(2):313–331, 2004.
- [24] Johan Håstad, Frank Thompson Leighton, and Brian Rogoff. Analysis of backoff protocols for multiple access channels. *SIAM Journal on Computing*, 25(4):740–774, 1996.
- [25] Elijah Hradovich, Marek Klonowski, and Dariusz R. Kowalski. New view on adversarial queueing on MAC. *IEEE Communication Letters*, 25(4):1144–1148, 2021.
- [26] Byung-Jae Kwak, Nah-Oak Song, and Leonard E. Miller. Performance analysis of exponential backoff. *IEEE/ACM Transactions on Networking*, 13(2):343–355, 2005.
- [27] Robert M. Metcalfe and David R. Boggs. Ethernet: Distributed packet switching for local computer networks. *Communications of the ACM*, 19(7):395–404, 1976.
- [28] Michael Mitzenmacher and Eli Upfal. *Probability and Computing*. Cambridge University Press, Second edition, 2017.
- [29] Adi Rosén and Michael S. Tsirkin. On delivery times in packet networks under adversarial traffic. *Theory of Computing Systems*, 39(6):805–827, 2006.
- [30] Qian M. Zhou, Aiden Calvert, and Maxwell Young. Singletons for simpletons: Revisiting windowed backoff with chernoff bounds. In *Proceedings of the 10th International Conference on Fun with Algorithms (FUN 2021)*, volume 157 of *LIPICs*, pages 24:1–24:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.