

Contention Resolution *without* Collision Detection*

Michael A. Bender
Dept of CS, Stony Brook University
Stony Brook, NY, USA
bender@cs.stonybrook.edu

William Kuszmaul
CSAIL, MIT
Cambridge, MA, USA
kuszmaul@mit.edu

Tsvi Kopelowitz
Dept. of CS, Bar-Ilan University
Ramat Gan, Israel
kopelot@gmail.com

Seth Pettie
Dept. of CS, University of Michigan
Ann Arbor, MI, USA
seth@pettie.net

ABSTRACT

This paper focuses on the contention resolution problem on a shared communication channel that does not support collision detection. A shared communication channel is a multiple access channel, which consists of a sequence of synchronized time slots. Players on the channel may attempt to broadcast a packet (message) in any time slot. A player's broadcast succeeds if no other player broadcasts during that slot. If two or more players broadcast in the same time slot, then the broadcasts collide and both broadcasts fail. The lack of collision detection means that a player monitoring the channel cannot differentiate between the case of two or more players broadcasting in the same slot (a collision) and zero players broadcasting. In the contention-resolution problem, players arrive on the channel over time, and each player has one packet to transmit. The goal is to coordinate the players so that each player is able to successfully transmit its packet within reasonable time. However, the players can only communicate via the shared channel by choosing to either broadcast or not. A contention-resolution protocol is measured in terms of its throughput (channel utilization). Previous work on contention resolution that achieved constant throughput assumed that either players could detect collisions, or the players' arrival pattern is generated by a memoryless (non-adversarial) process.

The foundational question answered by this paper is whether collision detection is a luxury or necessity when the objective is to achieve constant throughput. We show that even without collision detection, one can solve contention resolution, achieving constant throughput, with high probability.

CCS CONCEPTS

• Theory of computation → Parallel algorithms.

*This work was supported in part by NSF Grants CNS 1408695, CNS 1755615, CCF 1439084, CCF 1617618, CCF 1637546, CCF 1716252, CCF 1815316, XPS 1533644; by ISF Grant 1629/19 and BSF grant 2018364; and by an NSF GRFP Grant and a Fannie & John Hertz Foundation Fellowship.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

STOC '20, June 22–26, 2020, Chicago, IL, USA

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-6979-4/20/06...\$15.00

<https://doi.org/10.1145/3357713.3384305>

KEYWORDS

backoff, throughput, parallelism, networks

ACM Reference Format:

Michael A. Bender, Tsvi Kopelowitz, William Kuszmaul, and Seth Pettie. 2020. Contention Resolution *without* Collision Detection. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing (STOC '20)*, June 22–26, 2020, Chicago, IL, USA. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3357713.3384305>

1 INTRODUCTION

In the abstract contention resolution problem, there are multiple players that need to coordinate temporary and exclusive access to a shared resource. In this paper we use the terminology of one particular application, namely of many players, each of which must successfully transmit a single *packet* on a shared multiple-access communications channel. Contention resolution schemes are applied to many fundamental tasks in computer science and engineering, such as wireless communications using the IEEE 802.11 family of standards [1], transactional memory [36], lock acquisition [44], email retransmission [13, 22], congestion control (e.g., TCP) [37, 40], and a variety of cloud computing applications [32, 42, 50].

The classic algorithm for dealing with contention is the exponential backoff protocol [39]. The idea behind exponential backoff is that when a player p has a packet to send on the channel, then p must keep attempting broadcasts until p has a successful transmission (in which no other players were broadcasting). If a broadcast by p fails (due to a collision), then p waits for a random amount of time, proportional to how long p has been in the system, and then p attempts another broadcast.

Multiple access channels and the contention-resolution problem. Formally, the shared communication channel is modeled as a *multiple access channel*, which consists of a sequence of synchronized time slots (sometimes also called steps). Players on the channel may attempt to broadcast a packet (message) in any time slot. A player's broadcast *succeeds* (successfully transmits) if no other player broadcasts during that slot. If two or more players broadcast in the same slot, then the broadcasts *collide* and fail.

Although the slots are synchronized among all the players, there is no notion of a global clock (i.e., the players do not share a common time). Moreover, players are anonymous, i.e., players do not have ids. Players can see when successes occur on the channel, allowing for them to, for example, reduce their broadcast frequency when they see a long time interval without any successes. Players cannot

distinguish slots in which a collision occurred from slots in which no broadcast attempts were made, however.

In the **contention-resolution problem**, players arrive on the channel dynamically over time. A player p has one packet that p needs to broadcast. Player p automatically leaves the system once p 's packet has been successfully transmitted. The goal is to coordinate the players so that each player is able to successfully transmit its packet within reasonable time. We are interested in adversarial arrivals, that is, arrivals determined by an adaptive adversary. The adversary is able to see which time slots contain successes, and which player has succeeded in each of those steps. The input to the contention-resolution problem can be either a finite stream of n players, or an infinite stream of players.

Metrics. The primary objective of contention resolution is to optimize the **implicit throughput** (sometimes simply referred to as the throughput) of the channel. A slot is **active** if at least one player is in the system during that slot. For an n -player stream, the implicit throughput is defined to be n divided by the total number of active slots. More generally, if n_t is the number of players that arrive by time t and s_t is the number of active slots until time t , then the **implicit throughput** at time t is $\frac{n_t}{s_t}$, and the goal is to show that for each time t , the throughput is lower bounded by some positive constant with high probability in n_t . In other words, the number of active steps should be no more than a constant factor larger than the number of player arrivals.

By considering the ratio between players and active steps (rather than the ratio between successes and active steps), the implicit throughput gives each player the leeway to each cause collisions in a constant number of active steps without substantially damaging the metric performance. Bounds on the implicit throughput immediately imply bounds on player successes, in the following sense: If the implicit throughput is some constant c at step t , and if the number of players that arrive by time t is less than $\epsilon t/c$, then at least one of the most recent ϵt steps must be inactive, meaning that any player that arrived prior to step $(1 - \epsilon)t$ has succeeded. In Corollary 1, we present more general results relating constant implicit throughput to player successes.

Subject to small implicit throughput, we also want to minimize the number of **broadcast attempts** (i.e., we want the players to be energy efficient). That is, for an n -player input, players should, on average, make at most polylog n attempts to broadcast.

Achieving constant throughput with few broadcast attempts.

In this paper, we show that constant throughput is achievable, with high probability in the number of participating players. Moreover, if n players arrive in the first t steps, then with high probability in n , the average number of broadcast attempts per player is $O(\log^2 n)$.

Existing protocols and collision detection. Exponential backoff cannot even come close to achieving constant throughput (see the landmark paper by [2] and Section 2). Indeed, as we explain in Section 2, it is not hard to construct a situation where an arbitrarily small constant fraction of steps contain player arrivals, but at most a $1/\text{poly}(n)$ fraction of the slots are successes [7].

The challenge with designing contention-resolution protocols on a multiple-access channel is how to interpret failures in slots. Are these failures due to collisions or silence? How can a player use

the information provided by the channel to adjust the broadcast probabilities? The players' behavior provably needs to be a function of what they hear on the channel. That is, the players need to adjust (increase or decrease) their broadcast probabilities based on observing the channel (which slots have successful transmissions). Until now it has been unknown how to use such a mechanism to achieve $\Theta(1)$ throughput.

Thus, previous claims on constant throughput rely on restricting the model of player arrivals [28, 29, 31, 35, 43] and/or altering the communication channel to convey extra information to the players—for example, by having a **collisions detector** [10, 12, 15, 16].

The idea of collision detection is that when a player listens on the channel, not only does the player hear successful transmissions, but also *the player can distinguish between slots that are empty and slots with collisions*.¹ Collision detection allows players to differentiate between periods of true quiescence and high contention, typically increasing broadcast probabilities in the former case and decreasing in the latter. Collision detection has been a critical tool in designing protocols, e.g., by allowing protocols to use “busy signals” as a synchronization mechanism [10] or a multiplicative-weight update approach to discovering optimal transmission probabilities [15]. Achieving constant throughput with collision detection is relatively straightforward using the notion of a busy signal²; it becomes more challenging to simultaneously achieve other properties, such as resistance to a jamming adversary [10, 15] or guaranteeing $O(1)$ expected transmission attempts [15] per player.

In order to eliminate the need for collision detection, researchers have considered weaker versions of the contention-resolution problem. De Marco and Stachowiak considered a setting in which, even after a player successfully transmits, it can continue to send messages within the system [24]. This allows for a single player to be elected as a leader, and then to send a large stream of messages (one message ever $O(1)$ steps) in order to help synchronize other players. In this setting the authors were able to achieve constant throughput³ for a finite stream of messages [23]. Garncarek et al. considered a related problem in which players never leave the system, and instead are given new messages to send [25]; the algorithms for this problem have required that players have access to a global clock, however. Whether constant throughput can be achieved for the vanilla contention-resolution problem, without the use of conflict detection, has remained an open question.

Our results. In this paper we prove that it is possible to solve contention resolution with constant throughput without collision detection. In particular, let n_t be the total number of players to arrive in the first t steps. We guarantee that with high probability in n_t , the protocol achieves throughput of $\Omega(1)$ in the first t steps.⁴

¹Traditional hardware does not support collision detection in a reasonable way [4].

²A **busy signal** occurs when some synchronized subset of players all broadcast on every other step in order to signify to all other players in the system that they should be silent; this allows for the synchronized subset of players to all achieve their successes without ever having to interact with newly arriving players.

³The definition of throughput used in [23] is slightly weaker than the one in this paper. Rather than considering the ratio of players to active slots, they consider the ratio of players to the longest lifespan of any player (i.e., the most consecutive active slots during which any individual player is active).

⁴Without loss of generality, $n_t \leq O(t)$. In particular, if $n_t \geq \Omega(t)$, then the number of player arrivals per active step will necessarily be $\Omega(1)$, regardless of the protocol.

Moreover, the players make a total of $O(n_t \log^2 n_t)$ transmission attempts in the first t steps.

We also investigate the contention-resolution problem in the presence of adversarial jamming (the adversary has the power to block the broadcast attempts in some slots). Here, we prove that collision detection is fundamentally necessary. Whereas algorithms using collision detection [10, 15] can handle a small fraction of slots being jammed while still maintaining high-probability guarantees on the throughput, we prove a lower bound prohibiting any algorithm without collision detection from achieving even a *first* broadcast in time linear in the number of players with better than constant probability in the number of players. Moreover, this lower bound holds even when the adversary determining the jammed slots and arrival times of players is oblivious.

Additional related work. Willard [52] considered a contention-resolution problem where the goal is to minimize the *first* time that any player transmits. Sharp time bounds of $\Theta(\log \log N)$ (in expectation) are proved when the N players begin at the same time.

For many years, most of the analytic results on backoff assumed statistical queuing-theory models and focused on the question of what packet-arrival rates are stable (see [28, 29, 31, 35, 43]). Interestingly, even with Poisson arrivals, there are better protocols than binary exponential backoff, such as polynomial backoff [35]. The notion of saturated throughput—roughly, the maximum throughput under stable packet arrival rates—has been examined [14, 51]. The guarantees in our paper are much stronger because we guarantee constant utilization for *arbitrarily large* arrival rates.

There has been work on adversarial queuing theory, looking at the worst-case performance of these protocols [3, 5, 7, 8, 19, 20, 27, 30, 33, 52]. A common theme throughout these papers, however, is that dynamic arrivals are hard to cope with. When all the players begin at the same time, efficient protocols are possible [5, 7, 8, 26, 27, 33, 34, 52]. When players begin at different times, the problem is harder. The dynamic-arrival setting has been explicitly studied in the context of the wake-up problem [17, 18, 21], which looks at how long it takes for a single transmission to succeed when packets arrive dynamically.

A number of elegant results exist on contention resolution when the channel is subject to (possibly malicious) noise [6, 41, 45–49]. A recent result [9, 15] also addresses worst-case online arrivals of players and, in the face of an unknown j noisy slots scheduled by an adaptive adversary, achieves expected constant utilization with an expected $\text{polylog}(n + j)$ number of broadcasts per player.

Paper outline. The remainder of the paper proceeds as follows. In Section 2 we give a technical overview of our algorithm and of the ideas behind its design. In Section 3 we describe the algorithm formally, and in Section 4 we give technical preliminaries needed for the analysis of the algorithm. In Sections 5 and 6, we develop lemmas bounding the time needed by protocols in order to achieve their *first* success. These lemmas then play a critical role in Section 7, where we prove our main result concerning the (implicit) throughput of our algorithm. Finally, in Sections 8 and 9, we analyze the energy efficiency of our algorithm (i.e., the average number of broadcast attempts by each player) and we prove a lower

bound prohibiting algorithms from behaving well (with better than constant probability) in the presence of a jamming adversary.

2 TECHNICAL OVERVIEW

We begin by discussing what causes exponential backoff to fail, and why batch protocols are easier.

Contention and exponential backoff. Define the contention C_t on a slot t to be the sum of the broadcast probabilities on that slot, i.e., the expected number of players that attempt a broadcast during that slot. In order for a slot t to have a constant probability of a successful transmission, C_t must be $\Theta(1)$. When $C_t = \Omega(1)$, the probability of a successful transmission in slot t is $2^{-\Theta(C_t)}$, and when $C_t = O(1)$ the probability of a successful transmission is $\Theta(C_t)$. Thus, if we want n players to run a contention-resolution protocol and achieve a constant throughput, then we need $\Theta(n)$ of the active slots to have constant contention.

To develop intuition, we use a contention-based argument to show that exponential backoff does not achieve $\Theta(1)$ throughput.

In exponential backoff, the probability that a player broadcasts in slot t (conditioned on no prior successful transmission) is $\Theta(1/t)$. This means that if a player is in the system for $\Theta(n)$ time steps, then that player makes $\Theta(\log n)$ broadcast attempts. Now suppose all n players spend $\Theta(n)$ time in the system (which an adversary can guarantee with a $\Theta(n)$ -sized burst of player arrivals). Then the sum of the contentions of all slots is $\Theta(n \log n)$. Achieving constant throughput means having $\Theta(n)$ active slots, which means an average of $\Theta(\log n)$ contention per slot. If an adversary spreads player arrivals over $\Theta(n)$ slots (after the initial burst of arrivals), then this contention is distributed so that every slot has $\Omega(\log n)$ contention. But then at most a $1/\text{poly}(n)$ fraction of the packets can have been successfully transmitted, and there cannot be constant throughput.

Why the batch problem is easier. In contrast, a synchronized *batch protocol*, where all the players arrive at the same time, *can* achieve constant throughput [7]. Because the players start synchronized, they can use repeated doubling to guess the value of n . Then as players succeed, the players progressively back on, increasing the broadcast probabilities as more players transmit successfully and leave the system. The success of a batch strategy is not a consequence of having small average contention. In fact, the batch strategy described above has an average contention of $\Theta(\log^2 n)$ [7]. However, because the players are synchronized, this contention is spread out *unevenly* so that a constant fraction of the slots do have constant contention.

Interestingly, even when the players arrive in a batch, exponential backoff does not achieve constant throughput due to the fact that it backs off in its broadcasting probabilities but does not back on as players succeed [7]. It is worth understanding how a batch instance of exponential backoff fails to achieve constant throughput, because our protocol uses this analysis to its advantage. When an n -player batch instance of exponential backoff starts, at first the contention is too high, and essentially no slots are successful. But about $\Theta(n)$ slots into the protocol, there are $\Theta(n)$ slots, all of which have $\Theta(1)$ contention. Then, the contention gradually drops to $o(1)$ with the result that stragglers stay in the system for $\Omega(n \text{ polylog}(n))$

time before they all succeed in transmitting [7]. Thus, for exponential backoff on batch arrivals, even though the protocol does not achieve constant throughput overall, a constant fraction of the first $\Theta(n)$ time steps are successful.

2.1 Components of Our Protocol

The players use the even-numbered slots and the odd-numbered slots to simulate two separate channels. Because the players cannot access a global clock, there is no global agreement as to which channel is odd-numbered versus even-numbered.

At any given moment, each player in our protocol performs a variant of exponential backoff on either one or both of the channels. The key algorithmic contribution of the protocol is a simple set of rules that allow for each player to decide on which channel(s) to perform exponential backoff (and with what backoff parameters) in a way that ensures high throughput.

Two alternating channels and an invariant. Channels are used to run synchronized batches, while maintaining the invariant that at any time, at most one of the channels is running a batch protocol, and the other channel is silent.

When a synchronized collection of players begin a batch protocol on one channel, they jam the other channel in order to keep it silent.

The jamming is performed probabilistically so that, once the batch protocol has run for time roughly proportional the number of players involved in it, the jamming on the silent channel will fail with high probability.

When a successful transmission occurs on the silent channel, the two channels then reverse roles, with the previously silent channel becoming the batch channel.

Selecting which channel *not* to be on. When a player first arrives in the system, the player has no way of knowing which channel is currently in batch-mode and which is currently silent. Ideally, the player would stay silent until seeing a successful transmission on one of the two channels, allowing them to identify which channel is currently in batch-mode. The player cannot do this, however, since it may be that the only players currently in the system are newly arrived players, and that no channel is currently performing a batch. Thus, newly arriving players select a channel arbitrarily and perform exponential backoff on that channel.

The newly arrived player continues to perform the exponential backoff until seeing at least one success on some channel, at which point the player knows to treat that channel as the current batch channel. Once the player has identified the current batch channel, the player queues (i.e., performs backoff) on the silent channel and waits to join the next batch protocol.

Exponential backoff: give up on constant throughput and aim for one success. When new players enter the system, the exponential backoffs that they perform can have the effect of essentially jamming the current batch-protocol channel.

Nonetheless, we show that the total time that a batch-protocol spends being jammed by new arrivals is, with high probability, at most proportional to the size of the batch plus the total number of new arrivals. The key insight is that, although the exponential backoffs performed by new arrivals do not obtain good throughput (and, can in fact severely diminish throughput), they are efficient

at achieving at least one (total) success in time proportional to the number of exponential backoffs being concurrently performed.

In particular, consider n new arrivals that all perform exponential backoffs on some channel (and with different arrival times). In order for there to be collisions in all $\Theta(n)$ of the next slots of the channel, the contention in each slot must be at least $c \log n$ for a sufficiently large c . A simple counting argument for exponential backoff shows that there are not enough broadcasts in total to achieve contention $c \log n$ in more than a constant fraction of the slots. Thus, with high probability, although most slots have collisions, at least one slot will contain a success. Moreover, this continues to be true even if other activity is occurring on the same channel (e.g., a batch operation), unless that other activity is so dense that almost all of the n slots would have contained collisions anyway.

Three phases of the protocol. We now describe the protocol from the perspective of a player arriving in the system.

Phase 1: Channel choosing. The player arbitrarily chooses a channel and runs a version of an exponential-backoff protocol. In particular, the player performs *c-backoff* for some large constant c , in which for every interval of the form $(c^\ell, c^{\ell+1}]$, the player selects c random steps during which to broadcast.

The player continues to run the backoff protocol up until it sees a successful slot on some channel. Once the player sees a successful slot, then the player switches to the channel on which the success *did not occur*, and starts phase 2. Note that the players running phase 1 are not synchronized. Moreover, the successful slot that the player sees could come from a player running a later phase.

Phase 2: Batch synchronization. The player runs another *c-backoff* protocol on its chosen channel, up until it sees a successful slot on that channel. Once the player sees a successful slot, then it starts phase 3 (again, on that channel).

Note that the players running phase 2 are not synchronized, and can even conflict with players running phase 1 on the same channel. However, all the players running phase 2 on the channel switch to phase 3 at the same time.

Phase 3: Batch execution. The player runs a third (modified and tuned) backoff protocol, with the purpose of achieving a constant fraction of successes during the protocol. At step t of the batch protocol, each player broadcasts with probability exactly $\frac{1}{t}$.

While the player runs the batch protocol on its channel, it also runs a jamming protocol on the other channel. The purpose of the jamming is to prevent a new batch from immediately starting on the other channel. A player running the jamming protocol does not jam in every slot. Rather t time steps into the jamming protocol, the player broadcasts (jams) on the other channel with probability $\Theta((\log t)/t)$. The batch protocol ends when there is a jamming failure, resulting in a successful transmission on the other channel. At this point, any remaining players in the batch immediately switch to the batch protocol starting on the other channel.

The parameter $\Theta((\log t)/t)$ simultaneously serves two purposes. First, because the players jam probabilistically, there are not too many additional broadcast attempts per player, which helps minimize the subsidiary metric. Second, this imperfect jamming allows the batch protocol to end quickly once its length is sufficiently

large in $\Omega(b)$, where b is the number of players participating in the protocol; this prevents the protocol from continuing into the regime where the broadcast probabilities of players are too low to result in consistent successes.^{5,6}

2.2 Analysis Overview

In order to analyze the implicit throughput of our algorithm, an essential idea is to charge the length of each component of the algorithm to the number of player arrivals and successes that occur during that portion of the algorithm. Roughly speaking, each batch operation will, with high probability in the number of participating players, either contain a large number of new player arrivals (proportional to the length of the batch operation) or a large number of player successes (proportional to the length of the batch operation).

Analyzing batch operations with low interference. The first step in the analysis is to show that, if a batch operation consists of n players, and fewer than εn new players arrive during the batch operation (for some constant $\varepsilon \in (0, 1)$), then the first $O(n)$ slots of the batch operation contain $\Theta(n)$ successful transmissions (with high probability in n).

Let I denote the interval of steps $n/2, n/2 + 1, \dots, n/2 + dn$ in the batch operation for some large constant d . Let $I_j \subseteq I$ be the interval of steps between the j -th and $(j + 1)$ -th successes in I . Call I_j **light** if the number of new players that arrive during I_j is significantly smaller than $|I_j|$, and **heavy** otherwise. We show that for all $t \in \mathbb{N}$, the probability of an interval I_j being both light and of length t or greater is at most $\frac{1}{\text{poly}(t)}$. This implies that,

$$\sum_{j \in [1, n/10]} \begin{cases} |I_j| & \text{if } I_j \text{ light} \\ 0 & \text{otherwise} \end{cases} < \frac{dn}{2},$$

with high probability in n .

Since the sum of the lengths of the heavy intervals I_j is necessarily $O(n)$ (recall that heavy intervals are densely filled with arrivals), it follows that

$$\sum_{j \in [1, n/10]} |I_j| < dn,$$

and thus that there are, with high probability in n , at least $n/10$ successes during the interval I .⁷

Analyzing batch operations that overstay their welcome.

During a batch operation \mathcal{B} with n participants, the silent channel is jammed probabilistically in order to try to ensure that, with high probability in n , the jamming fails at some point in the first $O(n)$ steps (but not within the first dn steps for the constant d used in

the analysis above). However, activity by other players in the system (in either phases 1 or 2) could potentially interfere with the termination of the batch (by preventing a successful transmission on the silent channel).

To handle this, we consider the **amortized length** of a batch \mathcal{B} , which is defined to be 0 if the sum of the number of successes plus the number of new arrivals during the batch is at least $\Omega(\ell)$, where ℓ is the true length of the batch; and to be the true batch length ℓ otherwise. In the former case, we consider the length of \mathcal{B} to be charged, in an amortized sense, to the successes and new arrivals during the batch. The nice property that amortized lengths satisfy is that, for all $t \in \mathbb{N}$, the probability of a given batch \mathcal{B} having amortized length t or greater is $\frac{1}{\text{poly}(t)}$.

A unified analysis of first successes. Many parts of the algorithm analysis require us to argue that, under certain conditions, the probability of there being a long window of silence is small. This is necessary both to bound the sum of the lengths of the light intervals during a batch operation, as well as to establish that batch operations do not have large amortized lengths.

In order to unify these analyses, we define the notion of a **balanced protocol** \mathcal{P} , in which many (possibly non-synchronized) players are all following back-off-like protocols simultaneously. We use a balls-in-bins style analysis to show that, as long as the average contention is not too large during a balanced protocol, then at least one success will occur with high probability.

The analysis of the first success in a balanced protocol splits into two cases. When the average contention of the protocol is very small, one can focus on a constant number of players and show that, with high probability, one of them achieves a successful transmission. On the other hand, when the average contention of the protocol is larger (but still not too large), then the expected number of total successes becomes polynomially large, and thus it suffices to prove a concentration bound on the number of successes. To do this, we take advantage of the fact that each time a player makes a randomized decision about when to broadcast, the decision affects the total number of successful transmissions by at most ± 1 ; this enables the use of McDiarmid's inequality in order to prove the desired concentration inequality.

Analyzing throughput and energy efficiency. To analyze the (implicit) throughput of the system, we combine the analysis over many batches in order to show that, with high probability, the combined sizes of the batches that achieve poor throughput can be amortized to the total number of player arrivals in the system.

To analyze energy efficiency, we wish to show that each player, on average, only makes $O(\log^2 n)$ total broadcast attempts in the first $O(n)$ active steps. Since each player participates in phases 1 and 2 at most once, these phases contribute at most $O(\log n)$ attempted broadcasts per player. A single player could potentially engage in many different batches (and thus many different instances of phase 3). Nonetheless amortizing the broadcast attempts in each batch either to the successes during that batch, or to the new arrivals that occur during that batch, we can prove a high-probability bound on the average number of broadcast attempts per player.

⁵Because the probabilistic jamming performed by the batch-participants naturally fails on its own, it may seem unnecessary for the players in phase 2 to also perform exponential backoff on the silent channel. The purpose of the phase-2 exponential backoff is primarily to handle the case in which, in fact, no batch operation is occurring on the other channel (or all of the participants in the batch operation have succeeded).

⁶There are technical reasons as to why having players in the batch ever increase their broadcast probability is problematic (in particular, it makes the potential interference between players in phase 1 and players in phase 3 much more problematic), and so by using a probabilistic jamming protocol, the batch ends naturally without players having to increase their broadcasting probabilities to make the continuation of the protocol effective.

⁷Note that additional care must also be taken to ensure that the *very first* success in I occurs within a reasonably small time frame.

3 PROTOCOL COMPONENTS

Let c, c_1, c_2 be large constants, with c_1 sufficiently large as a function of c , and c_2 sufficiently large as a function of c_1 . (One can think of c_1 as being the geometric mean between c_2 and c). The algorithm will be defined in terms of c and c_2 ; the role of c_1 will appear only in the algorithm analysis.

Let α_1, α_2 refer to the odd-numbered-steps and even-numbered-steps, also known as **channels**; for a given channel α , we use $\bar{\alpha}$ to refer to the other channel. When we say that a player p executes some protocol on channel α we mean that p skips over steps in $\bar{\alpha}$, and uses only steps in α for the protocol.

3.1 c -Backoff

We begin by introducing a simple exponential backoff strategy called c -backoff, which is used in phases 1 and 2. The goal of c -backoff is to guarantee a success within a reasonable number of steps, thereby providing a synchronization mechanism.

Definition 1. For integer $c \geq 2$, a player performs a **c -backoff protocol** starting at time-step t as follows: For each $\ell \in \mathbb{N}$ where $\ell \geq 1$, the player selects a broadcast set B_ℓ of c random (and not necessarily distinct) time steps in the range $R_\ell = (t + c^\ell, t + c^{\ell+1}]$. The player then broadcasts during all time-steps in B_ℓ for all ℓ .

Notice that in Definition 1 there is no stopping condition, even after a success. This is by design since in our algorithms there may be a subset of the players in the system executing one protocol (such as c -backoff), while additional players may be executing another protocol, and so the executions may be jamming each other. Thus, having exactly one broadcasting player from the players that are executing c -backoff at a given time step does not guarantee a true success for that player. As a result, it is useful for analysis to treat c -backoff as continuing indefinitely.

Lemma 1. If a player p performs c -backoff for τ time, then the number of times p broadcasts is at most $c \log_c \tau$.

PROOF. Until time $c^{\lceil \log_c \tau \rceil} \geq \tau$, p broadcasts c times for each $1 \leq \ell \leq \lceil \log_c \tau \rceil - 1$, for a total of $c(\lceil \log_c \tau \rceil - 1) \leq c \log_c \tau$ broadcasts. \square

Lemma 2. If a player p performs c -backoff starting at time t , then at time $t + \tau$ the probability that p broadcasts is at least $\frac{1}{\tau}$ and at most $\frac{2c}{\tau}$.

PROOF. Let $k = \lceil \log_c \tau \rceil$. Then at time τ the probability of a broadcast is $\frac{c}{c^k - c^{k-1}} = \frac{1}{c^{k-1} - c^{k-2}}$. For the lower bound, $\frac{1}{c^{k-1} - c^{k-2}} \geq \frac{1}{c^{k-1}} \geq \frac{1}{\tau}$. For the upper bound, notice that $c^{k-2} \leq c^{k-1}/2$, and so $\frac{1}{c^{k-1} - c^{k-2}} \leq \frac{2}{c^{k-1}} = \frac{2c}{c^k} \leq \frac{2c}{\tau}$. \square

3.2 The Batch and Jamming Protocols

We now introduce the batch protocol and the jamming protocol used by players in phase 3.

The goal of the batch protocol is to utilize synchronization between a batch of players (all starting at the same time), in order to enable a constant fraction of those players to succeed in a reasonable amount of time. The jamming protocol is simultaneously used on the other channel in order to prevent activity on that channel.

Definition 2. A player performs a **batch protocol** starting at time-step t as follows: during the i -th step of the protocol, the player broadcasts with probability $\frac{1}{i}$.

The goal of the jamming protocol is to prevent unsynchronized players in the system from disturbing an execution of a (synchronized) batched protocol.

Definition 3. A player performs a **c_2 -jamming protocol** starting at time-step t as follows: during the i -th step of the protocol, the player broadcasts with probability $\frac{c_2 \log i}{i}$.

Balanced executions. An important property of both the batch protocol and the jamming protocol is that both protocols are balanced in the sense that the contribution of each protocol to the contention is well controlled. The following definition of a balanced protocol captures this notion formally in a way that will be useful in our analysis; specifically we will show that balanced protocols interact with c -backoff protocols in a constructive manner when attempting to achieve a first success.

Definition 4. Let $d > 1$ be some constant. An execution of a protocol \mathcal{P} starting at time 0 is said to be **(d, τ) -balanced** if the following conditions hold:

- **(Monotone size requirement)** If at step s , there are m_s players executing \mathcal{P} in the system, then m_0, m_1, m_2, \dots is a (weakly) monotonically decreasing sequence.
- **(Monotone probability requirement)** If at step s , each player that is executing \mathcal{P} broadcasts with probability q_s , then q_0, q_1, q_2, \dots is a (weakly) monotonically decreasing sequence.
- **(τ -lower bound requirement)** for $s > \tau$, $m_s q_s \leq \frac{\log \tau}{d}$.
- **(τ -upper bound requirement)** for $s \leq d^6 \tau$, $m_s q_s \geq \frac{d \log \tau}{\tau}$.

3.3 The Main Protocol

We are now prepared to present the algorithm in detail. Upon arrival, each player enters the following three phases, continuing until the player succeeds:

- (1) **Channel-choosing phase:** Execute c -backoff on an arbitrary α_i until witnessing a success on some channel α .
- (2) **Batch-synchronization phase:** Execute c -backoff on channel $\bar{\alpha}$ until witnessing a success on channel $\bar{\alpha}$.
- (3) **Batch-execution phase:** Execute a batch protocol on channel $\bar{\alpha}$ and a jamming protocol on channel α , until a success occurs on channel α , in which case restart the Batch-execution phase in channel α (while jamming channel $\bar{\alpha}$).

Defining the start and end points of batch protocols. The notion of when a batch protocol terminates is, at least initially, ambiguous. For example, one might assume that a batch protocol terminates after all of the players of the batch have succeeded, or alternatively assume that a batch terminates when a success is heard on the other channel. In order to avoid these types of ambiguity we allow a batch operation on a channel α to extend beyond the successes of all its participants, as long as there are still other players in the system and no successes have yet occurred on $\bar{\alpha}$.

We also need a clear definition for the beginning and end of a batch operation. We consider the success marking the beginning

of a batch operation not to be part of the batch operation, and the success marking the end of a batch operation (if there is such a success) to be part of the batch operation.

Basic properties. The algorithm has several useful properties, each of which can be proven by induction:

- (1) **Property 1.** There can be at most one batch operation taking place at a time.
- (2) **Property 2.** Any player broadcasting on a channel $\bar{\alpha}$ in which a batch operation is occurring either (a) is participating in that batch operation; or (b) arrived during the batch operation, is engaged in Phase 1 (Signal Generation), and will vacate the channel $\bar{\alpha}$ upon seeing any successes.
- (3) **Property 3.** During a batch operation in α , every player not in the batch operation either arrived after the batch operation began, or began Phase 2 (Synchronization) in channel $\bar{\alpha}$ when the batch operation began.

When analyzing the algorithm, we will often assume Properties 1 and 2 implicitly (since they are used quite heavily). When using Property 3, we reference it directly.

4 PRELIMINARIES

Throughout the paper, we say that an event occurs with **high probability** in n if the probability of the event occurring is $1 - \frac{1}{\Omega(n^c)}$ for a constant c of our choice (depending on the constants used to define the event).

An essential ingredient several of our proofs will be the use of McDiarmid's Inequality.

THEOREM 1 (McDIARMID'S INEQUALITY [38]). *Let X_1, \dots, X_m be independent random variables over an arbitrary probability space. Let F be a function mapping X_1, \dots, X_m to \mathbb{R} , and suppose F satisfies,*

$$\sup_{x_1, x_2, \dots, x_n, \bar{x}_i} |F(x_1, x_2, \dots, x_{i-1}, x_i, x_{i+1}, \dots, x_n) - F(x_1, x_2, \dots, x_{i-1}, \bar{x}_i, x_{i+1}, \dots, x_n)| \leq c,$$

for some $c > 0$ and for all $1 \leq i \leq n$. That is, if $X_1, X_2, \dots, X_{i-1}, X_{i+1}, \dots, X_n$ are fixed, then the value of X_i can affect the value of $F(X_1, \dots, X_n)$ by at most c . Then for all $R > 0$,

$$\Pr[F(X_1, \dots, X_n) - \mathbb{E}[F(X_1, \dots, X_n)] \geq R] \leq e^{-2R^2/(c^2n)},$$

and

$$\Pr[F(X_1, \dots, X_n) - \mathbb{E}[F(X_1, \dots, X_n)] \leq -R] \leq e^{-2R^2/(c^2n)}.$$

We will also make extensive use of the following lemma, which is a consequence of the Azuma-Hoeffding inequality for supermartingales with bounded differences. The proof appears in the extended paper [11].

Lemma 3. *Suppose X_1, \dots, X_n are (dependent) random variables such that*

$$\mathbb{E}[X_i \mid X_1 = a_1, X_2 = a_2, \dots, X_{i-1} = a_{i-1}] \leq O(1),$$

for all values a_1, \dots, a_{i-1} of X_1, \dots, X_{i-1} , and for all i . Moreover, suppose that deterministically $|X_i| \leq O(n^{0.1})$ for each i . Then with

probability $1 - n^{-\omega(1)}$,

$$\sum_i X_i \leq O(n).$$

Finally, it will be useful to have the following lemma which examines the probability that the sum of independent zero-one random variables takes value either 0 or 1. The proof appears in the extended paper [11].

Lemma 4. *Let $X = X_1 + \dots + X_t$ be the sum of t independent 0-1 random variables. Suppose each X_i takes value 1 with probability $p_i \leq 1/2$. Then*

$$\Pr[X = 1] \geq \Omega\left(\min(\mathbb{E}[X], 1/2^{2\mathbb{E}[X]})\right) \quad \text{and} \quad \Pr[X = 0] \geq 1/2^{2\mathbb{E}[X]}.$$

5 ANALYZING c -BACKOFF

Suppose that a c -backoff protocol is the only protocol being executed in a system, and that the first player arrives at time 0. In this section it is proven that if n players arrive by step τ for a carefully chosen τ , then there are at least \sqrt{c} successes by time τ (with high probability), and more specifically, during a carefully chosen interval of steps that ends at step τ . The reason for proving that there are at least \sqrt{c} successes (as opposed to just 1) is that in later analysis this will ensure there is still at least one success even when c -backoff is being executed in congruence with other protocols during the same steps; see Lemma 7. Whereas in general we will allow the arrival-times of players to be determined by an adaptive adversary, in this section we consider only oblivious adversaries; this is allowable because all applications of the lemma will be concerned with generating only a single success (and prior to the first success, adaptive and oblivious adversaries are indistinguishable).

We begin by focusing on a sparse case in which the number of players that have entered the system by time τ is polynomially smaller than τ . This case is formally stated in Lemma 5. We then focus on a more general case where the number of players that have entered the system by time τ is $O(\tau)$. This case is formally stated in Lemma 6. Notice that the case in which the number of players in the system by time τ is $\Omega(\tau)$ with a sufficiently large constant will not concern us.

Lemma 5. *Let $0 \leq \varepsilon \leq 1/2$ be a constant and let $c \in \mathbb{N}$ be a sufficiently large constant (such that $\sqrt{c} \gg 1/\varepsilon$). Suppose n players performing c -backoff arrive into the system by time $\tau = c^{k+1}$ for some integer $k \geq 2$, where the first player arrives at time 0 and the arrival times of the other players are determined by an oblivious adversary. If $n = O(\tau^{1-\varepsilon})$ then with high probability in τ there are at least \sqrt{c} successes in the range $(\tau/c, \tau]$.*

PROOF. Let p be the first player, and so p arrives at time 0. By the definition of the protocol, at time $\tau/c + 1$ player p enters the k -th integer range $R_k = (c^k, c^{k+1}] = (\tau/c, \tau]$ of the c -backoff protocol. In order to prove the lemma, we will prove that p has at least \sqrt{c} successes within the range R_k , with high probability in τ .

Player p selects a set B_k of c steps from R_k during which to broadcast. (In fact, B_k is a multiset, because the steps are chosen

with replacement.) Let $S_1, \dots, S_{\sqrt{c}}$ be a partition of B_k such that for $1 \leq i \leq \sqrt{c}$, $|S_i| = \sqrt{c}$.

For a given S_i in the partition, a step $s \in R_k$ is said to be **bad** if either $s \in B_k \setminus S_i$ or if one of the other $n - 1$ players (excluding p) chooses to broadcast during step s . The number of bad steps due to $B_k \setminus S_i$ is at most $c - \sqrt{c}$. By Lemma 1, the number of bad steps due to any one of the $n - 1$ other players is $O(c \log_c \tau)$, and so the number of bad steps due to all of the $n - 1$ other players is at most $O(cn \log_c \tau)$. Thus, the total number of bad steps, for any arbitrary choices of $B_k \setminus S_i$ and of the broadcast-times made by the $n - 1$ other players, is at most $O(cn \log_c \tau)$.

Notice that $|R_k| = \Theta(\tau)$. Since each element $s \in S_i$ is a random step in R_k , the probability that s is a bad step is

$$O\left(\frac{cn \log_c \tau}{|R_k|}\right) = \tilde{O}\left(\frac{\tau^{1-\varepsilon}}{\tau}\right) = \tilde{O}\left(\frac{1}{\tau^\varepsilon}\right).$$

Let A_i be the event that no broadcast in S_i was successful, which happens if and only if all of the steps in S_i are bad. Since the elements in S_i are independent, $\Pr[A_i] = \tilde{O}\left(\frac{1}{\tau^{\varepsilon\sqrt{c}}}\right)$. By the union bound, the probability that in each S_i at least one broadcast was successful is $1 - \tilde{O}\left(\frac{\sqrt{c}}{\tau^{\varepsilon\sqrt{c}}}\right)$. For a large enough choice of c , the number of successes is at least \sqrt{c} with high probability in τ . \square

Lemma 6. *Let $c \in \mathbb{N}$ be a sufficiently large square constant. Suppose n players performing c -backoff arrive into the system by time $\tau = c^{k+2}$ for some integer $k \geq 2$, where the first player arrives at time 0 and the arrival times of the other players is determined by an oblivious adversary. If $n \leq c^k$, then with high probability in τ there are at least \sqrt{c} successful transmissions in the range $(\tau/c^2, \tau]$.*

PROOF. Let $\tau' = \tau/c = c^{k+1}$, and fix a constant $0 \leq \varepsilon \leq 1/2$. There are two cases to consider depending on the number of players n' that arrived by time τ' . If $n' = O((\tau')^{1-\varepsilon})$ then, by Lemma 5, for a large enough choice of c , there are \sqrt{c} successes in the range $(\tau'/c, \tau'] \subset (\tau/c^2, \tau]$ with high probability in τ' (and thus in τ).

Next, consider the case where $n' = \Omega((\tau')^{1-\varepsilon})$. For each step $s \in (\tau', \tau]$, let b_s be the number of players that broadcast at step s . Thus, the contention at step s is, by definition, $\mathbb{E}[b_s]$. Step s is said to be **light** if $\mathbb{E}[b_s] \leq \frac{1}{c} \log \tau$ and **heavy** otherwise. Since $n \leq c^k$, and since by Lemma 1, each one of the players broadcasts at most $c \log_c \tau$ times by time τ , the number of broadcasts by all players within time range $(\tau', \tau]$ is at most $cn \log_c \tau$, and so,

$$\sum_{s=\tau'+1}^{\tau} \mathbb{E}[b_s] = \mathbb{E}\left[\sum_{s=\tau'+1}^{\tau} b_s\right] \leq cn \log_c \tau.$$

Thus, the number of heavy steps $s \in (\tau', \tau]$ is at most $c^2 n / \log c$, which is at most a $\frac{c^2 n}{(\tau - \tau') \log c}$ fraction of the steps in $(\tau', \tau]$. Since $\tau' = \tau/c \leq \tau/2$ and since $c^2 n \leq c^{k+2} = \tau$, the fraction of light steps in $(\tau', \tau]$ is at least

$$1 - \frac{c^2 n}{(\tau - \tau') \log c} \geq 1 - \frac{2c^2 n}{\tau \log c} \geq 1 - \frac{2}{\log c}.$$

Thus, if $c \geq 16$ then at least half of the steps in $(\tau', \tau]$ are light.

On the other hand, since there are at least $n' = \Omega((\tau')^{1-\varepsilon})$ players in the system at time $s \in (\tau', \tau]$ and, by Lemma 2, each one of these players broadcasts at step s with probability of at least $\frac{1}{\tau - \tau'}$, then

$\mathbb{E}[b_s] \geq \frac{1}{\tau - \tau'} \Omega(\tau'^{1-\varepsilon}) = \Omega(\tau^{-\varepsilon})$. Therefore, for any light step $s \in (\tau', \tau]$ we have

$$\Omega(\tau^{-\varepsilon}) \leq \mathbb{E}[b_s] \leq \frac{1}{c} \log \tau.$$

Notice that b_s depends on the random choices made by the players that are in the system at time s , and, in particular, b_s is a sum of independent random 0-1 variables, each of which takes value 1 with probability at most $1/(c-1) \leq 1/2$ (assuming $c \geq 3$). Thus, by Lemma 4, the probability of a success at a light step s is at least

$$\Omega\left(\min(\mathbb{E}[b_s], 2^{-2\mathbb{E}[b_s]})\right) = \Omega\left(\min(\tau^{-\varepsilon}, \tau^{-2/c})\right).$$

By setting $c \geq \frac{2}{\varepsilon}$, the probability becomes $\Omega(\frac{1}{\tau^\varepsilon})$.

Let N be the number of successes in $(\tau', \tau]$. Since at least half of the steps in $(\tau', \tau]$ are light steps, $\mathbb{E}[N] = \Omega(\tau^{1-\varepsilon})$.

In order to obtain a high probability bound on N notice that N is a function of $O(n \log \tau)$ independent random variables (specifically, these correspond with each of the time steps chosen by each player during which to broadcast). Moreover, changing a single variable can affect N by at most 1. It follows by McDiarmid's inequality (Theorem 1), that with probability $1 - 1/\tau^{\omega(1)}$, N does not deviate from $\mathbb{E}[N]$ by more than $\tau^{0.4} = o(\tau^{1-\varepsilon})$ (since $\varepsilon \leq 1/2$). Thus with (very) high probability in τ , $N = \Omega(\tau^{1-\varepsilon})$ which is much larger than \sqrt{c} , as required. \square

6 FIRST SUCCESS IN THE PRESENCE OF BALANCED EXECUTIONS

Recall that the intuition described in the previous section for requiring at least \sqrt{c} successes in Lemma 6 was that if the c -backoff protocol is not the only protocol being executed in the system, then multiple successes may be necessary before one of them is able to avoid conflicting with other concurrent protocols. Lemma 7 generalizes Lemma 6 to the more advanced setting in which there are also players in the system that are not performing c -backoff.

Lemma 7. *Let $c \in \mathbb{N}$ be a sufficiently large square constant. Suppose n **primary players** performing c -backoff arrive into the system by time $\tau = c^{k+5}$ for some integer $k \geq 2$, where $n \leq c^k$, and where the arrival times of the primary players are determined by an oblivious adversary. Suppose additionally that there are **secondary players** in the system participating in a (c, c^k) -balanced execution of some protocol. Then with high probability in c^k there will be at least one success in the time interval $(c^k, c^{k+5}]$.*

PROOF. Throughout the proof, for a step s let m_s be the number of secondary players in the system at step s and let q_s be the probability of a secondary player broadcasting at step s . Due to the monotone size and monotone probability requirements, the sequences $M = m_0, m_1, m_2, \dots$ and $Q = q_0, q_1, q_2, \dots$ are (weakly) monotonically decreasing.

The proof begins by considering the case where there are only secondary players during the first c^{k+1} steps.

Claim 1. *If no primary players arrive during the first c^{k+1} steps then at least one of the steps in the interval $(c^k, c^{k+1}]$ is a success of a secondary player, with high probability in c^k .*

PROOF. By the c^k -upper bound requirement, for all $s \leq c^{k+1}$ we have that $m_s \cdot q_s \geq \frac{c \log c^k}{c^k}$. In this case, we show that with high probability in c^k there is a success for some $s \in (c^k, c^{k+1}]$. Since $s > c^k$, by the c^k -lower bound requirement we have $m_s \cdot q_s \leq \frac{\log c^k}{c}$. Since there are no primary players during $(c^k, c^{k+1}]$, only secondary players are broadcasting during $(c^k, c^{k+1}]$. Thus, by Lemma 4, for each $s \in (c^k, c^{k+1}]$, the probability of a success at step s is at least

$$\Omega \left(\min \left(\frac{c \log c^k}{c^k}, \frac{1}{2 \frac{2 \log c^k}{c}} \right) \right),$$

where the constant in the Ω is independent of c . Since c is a sufficiently large constant the probability of a successful broadcast at each such step s is $\Omega(\frac{\sqrt{c} \log c^k}{c^k})$ (where again the constant in the Ω is independent of c). Since the steps in the interval $(c^k, c^{k+1}]$ are independent, the probability that at least one such step successfully broadcasting is at least with high probability in c^k , at least

$$\left(1 - \Omega \left(\frac{\sqrt{c} \log c^k}{c^k} \right) \right)^{c^k/2} \leq \frac{1}{\text{poly } c^k}.$$

□

It remains to consider the more general case in which at least one primary player arrives during the first c^{k+1} steps. For simplicity, relabel the step at which the first primary player arrives as time 0. The rest of the proof shows that with high probability in c^k there is a success in relabeled interval $(c^k, c^{k+4}]$ (which in the original relabeling is a subinterval of interval $(c^k, c^{k+5}]$).

The rest of the proof focuses on two cases, depending on whether the contention contributed by the secondary players at relabeled time c^{k+2} is high (at least $\frac{1}{c^{0.1k}}$) or low (at most $\frac{1}{c^{0.1k}}$). The claims for both cases follow the natural intuition that if the contribution of the secondary players to the contention is low then one of the primary players will successfully broadcast, and if the contribution of the secondary players to the contention is high (but not too high due to the c^k -lower bound requirement) then one of the secondary players will successfully broadcast.

Claim 2. *If at least one primary player arrives at time 0 and $m_{c^{k+2}} \cdot q_{c^{k+2}} \leq \frac{1}{c^{0.1k}}$, then there will be at least one success by a primary player during interval $(c^{k+2}, c^{k+4}]$, with high probability in c^k .*

PROOF. By Lemma 6, with high probability in c^k , there exist at least \sqrt{c} steps in $(c^{k+2}, c^{k+4}]$, denoted by $s_1, s_2, s_3 \dots$, such that for each s_i there is either a success prior to s_i in $(c^{k+2}, c^{k+4}]$ or exactly one primary player will broadcast at step s_i . For the following, focus on the first \sqrt{c} of these steps.

For each $1 \leq i \leq \sqrt{c}$, the secondary players have probability at most $m_{s_i} \cdot q_{s_i} \leq \frac{1}{c^{0.1k}}$ of conflicting with s_i . Thus, the probability of the secondary players conflicting with all s_i for $1 \leq i \leq \sqrt{c}$ is at most $\frac{1}{c^{0.1kc}}$. If c is a large enough constant, then with high probability in c^k , there exists $1 \leq i \leq \sqrt{c}$ such that s_i is a success. □

Claim 3. *If at least one primary player arrives at time 0 and $m_{c^{k+2}} \cdot q_{c^{k+2}} > \frac{1}{c^{0.1k}}$, then there will be at least one success by a secondary player during the interval $(c^k, c^{k+1}]$.*

PROOF. By the c^k -lower bound requirement, for each step $s > c^k$, we have $m_s \cdot q_s \leq \frac{\log c^k}{c}$. Moreover, by assumption and the monotonicity of M and Q , for each $s \leq c^{k+2}$, we have $m_s \cdot q_s > \frac{1}{c^{0.1k}}$. Thus, for a sufficiently large c , by Lemma 4, for each $s \in (c^k, c^{k+2}]$, the probability that exactly one secondary player broadcasts is at least $\Omega(c^{-0.1k})$.

For each step $s \in (c^k, c^{k+1}]$, let b_s be the number of primary players that broadcast at step s . Notice that the contribution of primary players to the contention at step s is $\mathbb{E}[b_s]$. Step s is said to be **light** if $\mathbb{E}[b_s] \leq \frac{1}{c} \log c^k$ and **heavy** otherwise. Since $n \leq c^k$, and since by Lemma 1, each one of the primary players broadcasts at most $c \log_c c^{k+1} = c(k+1)$ times by time c^{k+1} , then the number of broadcasts by all primary players within time range $(c^k, c^{k+1}]$ is at most $c^k(k+1)$, and so,

$$\sum_{s=c^k+1}^{c^{k+1}} \mathbb{E}[b_s] = \mathbb{E} \left[\sum_{s=c^k+1}^{c^{k+1}} b_s \right] \leq c^k(k+1).$$

Thus, the number of heavy steps $s \in (c^k, c^{k+1}]$ is at most $c^{k+1}(k+1)/(k \log c)$, which is at most a $\frac{c^{k+1}(k+1)}{(c^{k+1}-c^k)k \log c}$ fraction of the steps in $(c^k, c^{k+1}]$. Since $c^k = c^{k+1}/c \leq c^{k+1}/2$, the fraction of light steps in $(c^k, c^{k+1}]$ is at least

$$1 - \frac{c^{k+1}(k+1)}{(c^{k+1}-c^k)k \log c} \geq 1 - \frac{2c^{k+1}(k+1)}{(c^{k+1})k \log c} = 1 - \frac{2(k+1)}{k \log c} > 1 - \frac{4}{\log c}.$$

Thus, if $c \geq 32$ then at least half of the steps in $(c^k, c^{k+1}]$ are light.

Notice that b_s depends on the random choices made by the primary players that are in the system at time s , and, in particular, b_s is a sum of independent random 0-1 variables, each of which takes value 1 with probability at most $1/(c-1) \leq 1/2$ for $c \geq 3$. Thus, for $c \geq 5$, by Lemma 4, every light step s satisfies,

$$\Pr[b_s = 0] = \Omega \left(\frac{1}{2^{2\mathbb{E}[b_s]}} \right) = \Omega \left(\frac{1}{2^{\frac{2}{c} \log c^k}} \right) = \Omega \left(\frac{1}{c^{0.4k}} \right).$$

Let N be the number of slots in $(c^k, c^{k+1}]$ in which no primary players are broadcasting. Since at least half of the steps in $(c^k, c^{k+1}]$ are light steps, $\mathbb{E}[N] = \Omega((c^{k+1}-c^k)/c^{0.4k}) = \Omega(c^{0.6k})$.

In order to obtain a high probability bound on N notice that N is a function of $O(c^k \log_c c^k)$ independent random variables, which are the decisions of when to broadcast for each of the at most c^k players. Moreover, changing a single variable can affect N by at most 1. It follows by McDiarmid's inequality (Theorem 1), that with probability $1 - 1/c^{\omega(k)}$, N does not deviate from $\mathbb{E}[N]$ by more than $c^{0.5k}$. Thus with (very) high probability in c^k , $N = \Omega(c^{0.6k})$, and so the probability of there being a step in which a single secondary player broadcasts and no primary player broadcasts is at least

$$1 - \left(1 - \Omega \left(\frac{1}{c^{0.1k}} \right) \right)^{\Omega(c^{0.6k})},$$

which is high probability in c^k . □

Thus the lemma is proven. □

Remark 1. *Lemmas 6 and 7 treat the arrival times and broadcast probabilities of players as being predetermined by an oblivious adversary. Recall that, in general, however, we wish to consider adaptive*

adversaries against our algorithm. Lemmas 6 and 7 will only be applied to settings in which we are attempting to obtain a first success, however, and in these settings an adaptive adversary has no additional power over an oblivious adversary (since there are no successes to adapt to). One subtlety, however, is that when we apply Lemmas 6 and 7, we will typically be applying them to only one channel α (i.e., the odd-indexed steps or the even-indexed steps), while other activity occurs on the other channel $\bar{\alpha}$. Critically, the activity on channel $\bar{\alpha}$ will only affect when (a) new backoff players arrive on channel $\bar{\alpha}$ and (b) the values of m_0, m_1, \dots and q_0, q_1, \dots ; and the randomness used by players in channel $\bar{\alpha}$ will be independent of the activity (and the random bits used) in channel α (at least until after the next success in channel α). Thus one can think of the arrival times of players on channel α , and the values of m_0, m_1, \dots and q_0, q_1, \dots as being fully determined by $\bar{\alpha}$ prior to the application of Lemma 6 or 7 to α .

Remark 2. For the sake of avoiding clutter, throughout the rest of the paper we use Lemma 7 without forcing τ to be a power of c . This relaxation adds at most a factor of c to the length of the interval in which there is at least one success with high probability.

7 ANALYZING THROUGHPUT

In this section, we analyze the (implicit) throughput of our algorithm, using Lemma 7 as an important building block. We will prove the following theorem about implicit throughput:

THEOREM 2. Recall that a step is active if at least one player is present during that step. Suppose n players arrive in the first t time steps. Then with high probability in n , at most $O(n)$ of the first t time steps are active.

We will also prove a corollary transforming the implicit-throughput result of Theorem 2 into a statement about the success times of players.

Corollary 1. Call a time-step t k -smooth if for all $j \geq k$, the number of arrivals in steps $t - j + 1, \dots, t$ is sufficiently small in $O(j)$. If a time-step t is k -smooth, then with high probability in k , all players that entered the system prior to step $t - k + 1$ are no longer in the system after time-step t .

We begin the proof of Theorem 2 by considering the number of successes within the first $O(n)$ steps of an n -player batch operation \mathcal{B} , assuming that not too many new players arrive during the execution of the batch operation. In particular, we will use Lemma 7 to show that, even though non-batch-operation players can add substantial contention to a given time slot of \mathcal{B} , at least a constant fraction of the time slots in \mathcal{B} will, with high probability in n , contain only batch-operation players; these steps will then guarantee a large number of successes for the batch operation.

Lemma 8. Consider a batch operation \mathcal{B} involving n participants in channel α , and condition on at most n/c_1 players joining the system during the first $\tau = c_1 n$ steps of the batch operation (or the first $|\mathcal{B}|$ steps if $|\mathcal{B}| \leq c_1 n$). Then with high probability in n , there are $\Omega(n)$ successes during the first $\min(|\mathcal{B}|, \tau)$ steps of the batch operation.

PROOF. We begin by showing that, with high probability in n , either there are at least $n/2$ successes during the execution of \mathcal{B} , or $|\mathcal{B}| \geq \tau$. The proof is based on the intuition that if more than $n/2$

players are executing \mathcal{B} on channel α , then these players are, with high probability in n , preventing any success from channel $\bar{\alpha}$.

Claim 4. Assuming that c_2 is sufficiently large with respect to c_1 , then with high probability in n , there are either at least $n/2$ successes during \mathcal{B} , or \mathcal{B} lasts for at least τ steps.

PROOF. If there are fewer than $n/2$ successes (in channel α) during the first $\min(|\mathcal{B}|, \tau)$ steps of \mathcal{B} in channel α , then during each one of those steps in channel $\bar{\alpha}$ there will be at least $n/2$ players each broadcasting with probability at least $\frac{c_2 \log \tau}{\tau} = \frac{c_2 \log(c_1 n)}{c_1 n}$. For c_2 sufficiently large relative to c_1 , this ensures for each step in $\bar{\alpha}$, that with high probability in n , at least two players broadcast. Thus with high probability in n , either there are more than $n/2$ successes during the first $\min(|\mathcal{B}|, \tau)$ steps of \mathcal{B} , or there are no successes in $\bar{\alpha}$ during those steps (in which case $|\mathcal{B}| > \tau$). \square

Define $s_0 = 0$, and for $i > 0$, define s_i to be the step-number (counting only steps in α) of the i -th success during \mathcal{B} (or to be $|\mathcal{B}|$ if there is no such success). If $s_i > \tau$, then we truncate s_i to be τ (i.e., we cap each s_i by τ). For $i \geq 1$, let X_i be the interval $(s_{i-1}, s_i]$.

We prove that, with high probability in n ,

$$\sum_{i=1}^{n/10} |X_i| < \tau. \quad (1)$$

Thus, $s_{n/10} < \tau$ which together with Claim 4 implies that, with high probability in n , either there are at least $n/2$ successes in the first $\min(|\mathcal{B}|, \tau)$ steps of \mathcal{B} , or $s_{n/10} < \tau \leq |\mathcal{B}|$. In both cases, the number of successes in the first $\min(|\mathcal{B}|, \tau)$ steps of \mathcal{B} is at least $n/10$, as desired.

Call an interval $I = [i_1, i_2]$ **light** if the number of new players that arrive during I is at most $2|I|/c_1$, and **heavy** otherwise. Since at most n/c_1 players join the system during the first τ steps, the sum of the lengths of the heavy intervals X_i satisfies

$$\sum_{\text{heavy } X_i} |X_i| \leq n/2 \leq \tau/2.$$

To prove Eq. 1, it therefore suffices to show that

$$\sum_{i \in [1, n/10], X_i \text{ light}} |X_i| < \tau/2. \quad (2)$$

To prove Eq. 2, we begin by considering intervals X_i where $s_{i-1} \geq n/2$. The next claim establishes that the interval X_i is with high probability either heavy (in which case it does not contribute to Eq. 2) or is relatively small.

Claim 5. Consider arbitrary fixed values for s_1, \dots, s_{i-1} such that $s_{i-1} \geq n/2$, and such that $i \leq n/10$. Then, conditioning on s_1, \dots, s_{i-1} , and for any value $t \in \mathbb{N}$, we have that, with high probability in t , either $|X_i| \leq t$ or X_i is heavy.

PROOF. Without loss of generality, assume that t is at least a sufficiently large constant, since otherwise the claim trivially holds.

The value of s_i (and hence also of $|X_i|$) is determined by the first success following s_{i-1} (or by the termination of \mathcal{B}). In order for X_i to be light and to also satisfy $|X_i| > t$, there must be some $j \geq t$ (specifically $j = |X_i| - 1$) such that the interval $(s_{i-1}, s_{i-1} + j + 1]$ in α is light; the interval $(s_{i-1}, s_{i-1} + j]$ contains no successes; and $s_{i-1} + j + 1$ is at most τ (since s_i is defined to be truncated to at

most τ). Define j^* to be the smallest $j \geq t$ such that the interval $(s_{i-1}, s_{i-1} + j + 1]$ is light.⁸ Since $j^* \leq j$, it must also be that the interval $I_{j^*} = (s_{i-1}, s_{i-1} + j^*]$ contains no successes. To prove the claim, it suffices to show that the probability of I_{j^*} containing no successes is polynomially small in j^* (and thus also polynomially small in t).

The rest of the proof establishes that the execution of \mathcal{B} is $(c, n/2)$ -balanced (in α), thereby enabling an application of Lemma 7 on I_{j^*} . Let m_s denote the number of players executing \mathcal{B} that are still in the system at time s . Let q_s denote the probability that a player executing \mathcal{B} broadcasts at step s . For all $s \in X_i$, we have $q_s = \frac{1}{s}$ and $m_s = n - i + 1$, exactly. Thus, \mathcal{B} has both the monotone size requirement and the monotone probability requirement

Since $I_{j^*} \cup \{s_{i-1} + j^* + 1\}$ is light, then there are at most $2(j^* + 1)/c_1$ new players that enter the system during I_{j^*} . Recall that there are at least $9n/10$ players from \mathcal{B} that are still in the system at step s_{i-1} (since $i \leq n/10$), and so for any $s \in I_{j^*}$, we have $9n/10 \leq m_s \leq n$. Moreover, the probability that a player executing \mathcal{B} broadcasts during step s is $q_s = \frac{1}{s}$, and so $m_s \cdot q_s \in [\frac{9}{10} \cdot n/s, n/s]$. For any $s \in I_{j^*}$, since $s_{i-1} \geq n/2$ it must be that $s > n/2$. Moreover, since $s_i \leq \tau$ it must be that $s \leq \tau - 1 = c_1 n - 1$. Thus, for each $s \in I_{j^*}$, $\frac{9}{10 \cdot c_1} < m_s \cdot q_s \leq 2$, and so \mathcal{B} fulfills both the $n/2$ -lower bound requirement and the $n/2$ -upper bound requirement. Thus, the execution of \mathcal{B} is $(c, n/2)$ -balanced.

We now apply Lemma 7 where the secondary players are the players participating in \mathcal{B} during I_{j^*} , and the primary players are players that are executing c -backoff on channel α while \mathcal{B} is being executed during I_{j^*} (these players are a subset of the players that joined the system during I_{j^*}). Notice that the application of Lemma 7 is possible since the number of primary players is at most n/c_1 which is at most $n/2$ for $c_1 \geq 2$. Thus, as long as t is at least a sufficiently large constant, it follows by Lemma 7 that the probability of not having a success during I_{j^*} is at most $1/\text{poly}(j^*) \leq 1/\text{poly}(t)$. \square

By Claim 5 and Lemma 3,

$$\sum_{i \in [1, n/10], X_i \text{ light}, s_{i-1} \geq n/2} |X_i| < O(n),$$

with high probability in n . For c_1 sufficiently large, it follows that with high probability in n ,

$$\sum_{i \in [1, n/10], X_i \text{ light}, s_{i-1} \geq n/2} |X_i| < \tau/4.$$

To prove Eq. 2, it therefore suffices to show that with high probability in n ,

$$\sum_{i \in [1, n/10], X_i \text{ light}, s_{i-1} < n/2} |X_i| < \tau/4. \quad (3)$$

Notice that $\sum_{\text{light } X_i \subseteq [0, n/2]} |X_i|$ is trivially bounded by $n/2 \leq \tau/8$, assuming that $c_1 \geq 4$. Thus the only interval that it remains to consider is the X_i for which $s_{i-1} < n/2$ but $s_i > n/2$ (if such an

⁸Note that j^* is a function of when new players arrive after step s_{i-1} . Thus j^* is determined by the adaptive adversary who selects player-arrival times. However, since our analysis is only concerned with the value that j^* would take in the event that no successes were to happen in the interval $(s_{i-1}, s_{i-1} + j^*]$, we can think of j^* as being a function of only s_1, \dots, s_{i-1} (and possibly of random bits used by the adversary). Importantly this means that j^* can be thought of as being determined prior to the execution of steps $s_{i-1} + 1, s_{i-1} + 2, \dots$, rather than being a random variable depending on what occurs in those steps.

interval X_i exists). Notice that if $|X_i| \leq \tau/8$, then the proof will be complete. Assuming that c_1 is sufficiently large, the desired statement $|X_i| \leq \tau/8$ is implied by the following claim.

Claim 6. *With high probability in n , there is at least one success in the interval $[n/2, c^6 \cdot n/2]$.*⁹

PROOF. Note that the interval $[1, c^6 \cdot n/2]$ is necessarily light by virtue of the fact that at most n/c_1 players can arrive in the first τ steps of \mathcal{B} . Let m_s denote the number of players executing \mathcal{B} that are still in the system at time s . Let q_s denote the probability that a player executing \mathcal{B} broadcasts at step s . For any $s \in [n/2, c^6 \cdot n/2]$ (prior to the first success in the interval) we have $q_s = \frac{1}{s}$ and $n/2 \leq m_s \leq n$. Thus, \mathcal{B} has both the monotone size requirement and the monotone probability requirement. Moreover, $m_s \cdot q_s$ is in the range $[1/(2 \cdot c^6), 2]$. Assuming n is at least a sufficiently large constant (which is w.l.o.g. since otherwise the entire lemma is immediate), both the $n/2$ -upper bound requirement and the $n/2$ -lower bound requirement hold for \mathcal{B} .

We apply Lemma 7, where the secondary players are the players participating in \mathcal{B} , and the primary players are players that are executing c -backoff on channel α while \mathcal{B} is being executed (these players are a subset of the players that joined the system during \mathcal{B}). Notice that the application of Lemma 7 is possible since the number of primary players is at most n/c_1 which is at most $n/2$ for $c_1 \geq 2$. Thus, by Lemma 7, with high probability in n , there is at least one success in the interval $[n/2, c^6 \cdot n/2]$, as desired. \square

This completes the proof of the lemma. \square

Lemma 8 ensures that the first $O(n)$ steps of an n -player batch operation \mathcal{B} will either contain a large number of successful broadcasts, or a large number of new player arrivals. The length of the batch operation \mathcal{B} could be substantially larger than $O(n)$, however, which would damage the implicit throughput of the system. The next lemma shows that either the batch operation \mathcal{B} will be small (at most length $O(n)$), or the length of \mathcal{B} can be attributed to a large number of player arrivals during \mathcal{B} (which we will later use to amortize any damage incurred on the implicit throughput by \mathcal{B}).

Lemma 9. *Consider a batch operation \mathcal{B} involving n participants in channel α . Let l be the length of \mathcal{B} , which is the number of steps of \mathcal{B} in channel α . Define the **truncated length** \bar{l} of \mathcal{B} to be zero if either the number of successes or the number of new arrivals during \mathcal{B} is $\Omega(l)$, and to be l otherwise. Then, for all $t \in \mathbb{N}$,*

$$\Pr[\bar{l} = t] \leq \frac{1}{\text{poly}(t)}.$$

PROOF. We begin by considering the case in which $t \leq O(n)$. By Lemma 8, with high probability in n , either the number of successes or the number of arrivals during the first $\min(|\mathcal{B}|, c_1 \cdot n)$ steps is $\Omega(n)$. Therefore, with high probability in n (which is also high probability in t), \bar{l} is zero.

Next, consider the case in which t is sufficiently large in $O(n)$. Claim 7 shows that with high probability there are only three cases:

⁹See Remark 2 for an explanation regarding why the exponent here is 6 and not 5.

Claim 7. Consider a batch operation \mathcal{B} that starts at time 0 in channel α , and has n participants. If t is sufficiently large in $\Omega(n)$ then during the first t/c_1 steps on channel α , with high probability in t , one of the following occurs:

- \mathcal{B} terminates.
- All the participants in \mathcal{B} successfully broadcast.
- There are $\Omega(t)$ arrivals of new players.

PROOF. Notice that Property 3 of the algorithm ensures that all players performing c -backoff on channel $\bar{\alpha}$ began their c -backoff after (or at the same time as) the beginning of the batch \mathcal{B} . Let m_s denote the number of players participating in \mathcal{B} during the s -th step on channel $\bar{\alpha}$. Let q_s denote the probability of a player in \mathcal{B} to broadcast during the s -th step on channel $\bar{\alpha}$. Notice that the sequences m_0, m_1, m_2, \dots and q_0, q_1, q_2, \dots fulfill the monotone success requirement and the monotone probability requirement.

Suppose that \mathcal{B} does not terminate by step t/c_1 and that there are fewer than $O(n)$ arrivals of new players by step t/c_1 , since otherwise the lemma holds automatically. If $m_{t/c_1} = 0$ then the proof is complete. Thus, suppose that $m_{t/c_1} > 0$. Let $\tau = t/(c^6 c_1)$, and notice that $\tau = \Omega(n)$ (i.e., $\tau \geq c'n$ for a constant c' of our choice). Then, for step $\tau \leq s \leq t/c_1$, $\frac{c_2 c_1 \log(t/c_1)}{t} \leq q_s \leq \frac{c_2 \log \tau}{\tau}$ and $1 \leq m_s \leq n$. Thus, for a sufficiently large constant hidden in the $\Omega(n)$ bound on τ , $m_s q_s \in [\frac{c \log \tau}{\tau}, \frac{\log \tau}{c}]$. Thus, both the τ -upper bound requirement and the τ -lower bound requirement hold for \mathcal{B} on channel $\bar{\alpha}$, and so the execution of \mathcal{B} over channel $\bar{\alpha}$ is (c, τ) -balanced. Since the number of new players that arrived in channel $\bar{\alpha}$ is $O(n)$ and t is assumed to be sufficiently large in $\Omega(n)$, by Lemma 7 (with high probability in n) there is a success in channel $\bar{\alpha}$ during the first t/c_1 steps, thereby terminating \mathcal{B} , which is a contradiction. \square

In the first and third cases of Claim 7, we have that $\bar{l} < t$. Thus we need only focus on the second case, in which all of the players of \mathcal{B} successfully broadcast by step t/c_1 .

If all of the players of \mathcal{B} successfully broadcast by step t/c_1 and there are no more players in the system, then $\bar{l} < t$. Thus, the remaining option to consider is that there are still players in the system at step t/c_1 , but none of them are participating in \mathcal{B} . In this case, by Lemma 6, with high probability in t , there is at least one success in some channel between step t/c_1 and step $t/\sqrt{c_1}$. If the success takes place in channel $\bar{\alpha}$, then \mathcal{B} terminates and thus $\bar{l} < t$. To consider the remaining case, suppose that the success takes place in channel α and let r be the number of players in the system at the end of the step. If $r = 0$, then \mathcal{B} terminates (and thus again we have $\bar{l} < t$). Otherwise, if $r > 0$, then right after the success all of the remaining $r - 1$ players are running c -backoff in channel $\bar{\alpha}$. In this case, by applying Lemma 6, with high probability in t , there is a success in channel $\bar{\alpha}$ prior to step t , and again $\bar{l} < t$. In every case we get that with high probability in t , $\bar{l} \neq t$. \square

So far we have focused on the behavior of batch operations \mathcal{B} , showing that with good probability, the length of \mathcal{B} can be charged either to a proportionally large number of successes or to a proportionally large number of new player arrivals. The next lemma focuses on the active steps during which no batch operation is occurring. We show that with good probability, the number of

consecutive such steps prior to the start of the next batch operation (or prior to a non-active step) is either small or can be attributed to a large number of player arrivals during those steps.

Lemma 10. Consider an arrival at step t_0 such that all of the players that are in the system at time t_0 were not in the system at step $t_0 - 1$. Let s be the minimum between the number of steps before the beginning of the next synchronized batch operation, and the number of steps before the next time τ where all of the players in the system at time τ were not in the system at time $\tau - 1$. Let \bar{s} be 0 if there are $\Omega(s)$ arrivals during those s steps, or s otherwise. Then for all $t > 0$,

$$\Pr[\bar{s} = t + 1] \leq \frac{1}{\text{poly}(t)}.$$

PROOF. If there are $\Omega(t)$ arrivals during the first t steps, then $0 = \bar{s} \neq t$. Thus, the rest of the proof is conditioned on there being a sufficiently small $O(t)$ number of arrivals during the first t steps.

Let α be the channel on which the first player to arrive at time t_0 performs their c -backoff in Phase 1 of the algorithm. Then either there will be a success in channel $\bar{\alpha}$ during the first $t/2$ steps after t_0 , or, by Lemma 6, with high probability in t , there will be at least one success in channel α during the first $t/2$ steps after t_0 . Let β be the channel that achieves the first success after time t_0 , and condition on that success occurring within the first $t/2$ steps after t_0 . When the first success occurs, all players currently in the system will begin Phase 2 of the algorithm on channel $\bar{\beta}$. If there are no such players, then the system must be empty and $s = t/2 < t + 1$. Otherwise, apply Lemma 6 to deduce that, with high probability in t , a success occurs in channel $\bar{\beta}$ within the following $t/2$ time steps. Thus, with high probability in t , the first t steps will contain a success, marking the beginning of a batch operation, and so $\bar{s} < t + 1$, as desired. \square

The preceding lemmas analyze individual components of the algorithm in order to show that when a step in the algorithm takes a large number of steps, the number of such steps can be attributed either to a correspondingly large number of successes or to a correspondingly large number of new player arrivals. To complete the proof of Theorem 2, we use the Azuma-Hoeffding inequality (i.e., Lemma 3) in order to consider all the steps of the algorithm concurrently (rather than each individually), and then we perform an amortization argument in which we charge (almost all) active steps either to successes or to player arrivals, thereby bounding the number of active steps by $O(n)$, the number of player arrivals.

PROOF OF THEOREM 2. Let l_1, \dots, l_n be the lengths of the first batch operations that take place prior to step t (with $l_i = 0$ if fewer than i such batch operation occur), and let $\bar{l}_1, \dots, \bar{l}_n$ be the truncated lengths of the batch operations (as defined by Lemma 9). Similarly, let s_1, \dots, s_n be the lengths of the runs of active steps prior to step t in which no batch operation is occurring (with $s_i = 0$ if fewer than i such runs occur), and let $\bar{s}_1, \dots, \bar{s}_n$ be the truncated lengths of the runs (as defined by Lemma 10). By Lemmas 9, 10, and 3, we have with high probability in n that

$$\sum_{i=1}^n \bar{l}_i + \sum_{i=1}^n \bar{s}_i \leq O(n).$$

Since at most n arrivals and at most n successes occur during the first t steps, the definitions of the truncated values \bar{s}_i and \bar{l}_i imply that

$$\sum_{i=1}^n l_i + \sum_{i=1}^n s_i \leq \sum_{i=1}^n \bar{l}_i + \sum_{i=1}^n \bar{s}_i + O(n) \leq O(n).$$

Since the left-hand side is the number of active steps, the theorem follows. \square

We conclude the section by proving Corollary 1.

PROOF. Consider the event E that some player p that entered the system prior to step $t - k + 1$ is still present in the system after time-step t . Then all of steps $t - k + 1, \dots, t$ must be active. Therefore, if event E holds, then there exists some $j \geq k$ such that all of the j steps $t - j + 1, \dots, t$ preceding step t are active, and such that at the end of step $t - j$ no players were in the system. (Note that j may equal t if all steps prior to t have been active.) However, by Theorem 2 the probability of any particular j having this property is $\frac{1}{\text{poly}(j)}$. Summing over all $j \geq k$, we get that with high probability in k , the event E does not occur. \square

8 ANALYZING ENERGY EFFICIENCY

The next theorem analyzes the energy efficiency of the algorithm, bounding the total number of broadcast attempts by players to be, on average, at most polylogarithmic.

THEOREM 3. *Suppose n players arrive in the first t steps. Then there are at most $O(n \cdot \log^2 n)$ broadcast attempts in the first t steps, with high probability in n .*

We defer the proof of Theorem 3 to the extended paper [11].

9 THE POWER OF JAMMING

In this section we consider contention resolution in the presence of a **jamming adversary**, which is able to **jam a step** by forcing it to unconditionally fail at successfully transmitting.

We show that, in the presence of jamming, even achieving a *single* success becomes difficult. This demonstrates a fundamental separation between contention-resolution with and without collision detection. In particular, algorithms in which players are able to perform collision detection can not only achieve a first success, but can also achieve high implicit throughput even in the presence of adversarial jamming [15].

We begin by defining what it means for an algorithm whose goal is to achieve a single success to be resistant to jamming:

Definition 5. *Consider a contention resolution algorithm \mathcal{A} and let c be a sufficiently large constant. Consider a sequence of steps such that one player arrives at time 0; at most n/c additional players arrive before time n ; all arriving players perform algorithm \mathcal{A} ; and at most n/c steps are jammed by an adversary. We say that \mathcal{A} is **$f(n)$ -resilient to jamming** if in this setting, at least one success occurs in the first n steps with probability at least $1 - f(n)$ (and for any value of n).*

The next theorem establishes that no contention resolution algorithm \mathcal{A} can handle jamming with better than constant probability.

THEOREM 4. *No algorithm can be $f(n)$ -resilient to jamming for any monotone decreasing $f(n)$ satisfying $f(n) \leq o(1)$.*

PROOF. Suppose \mathcal{A} is $f(n)$ -resilient to jamming for some monotone decreasing function $f(n)$ satisfying $f(n) \leq o(1)$.

Consider an adversary that first jams the system for $\frac{n}{2c}$ steps (for convenience assume n is a power of $2c$), and then selects at random (with replacement) $\frac{n}{2c}$ additional steps in $[n]$ to jam. With probability at least $1 - f(n)$, \mathcal{A} must achieve at least one success against this adversary, even if there is only a single player p in the system (who arrives at time one). Define x to be a value sufficiently small in $O(\log_{1/c} f(n))$. If the player p attempts to broadcast fewer than x times in the interval $(n/2c, n]$, then p has a greater than $2 \cdot f(n)$ probability of failing to successfully transmit; thus \mathcal{A} must guarantee that with probability at least $1/2$, p broadcasts at least x times during the interval $(n/2c, n]$. Hence the expected number of broadcasts by p in interval $(n/2c, n]$ is $\Omega(\log \frac{1}{f(n)})$.

Summing over the intervals $(1, 2c], (2c, 4c^2], \dots, (n/2c, n]$, we get that the expected number of broadcasts by a player in its first n steps is asymptotically larger than $\log n$ (i.e., $\omega(\log n)$). We will now exploit this property, in order to reach a contradiction.

Consider a new adversary that jams only the steps $1, \dots, n/c$, inserts one player into the system prior to step one, and then inserts $n/c - 1$ players each at an independent random time step in $1, \dots, n$. For any step $t \in (n/c, n]$, and for a player p that is inserted at a random time step, we claim that the probability that p broadcasts at time t is at least $\omega(\log n/n)$. In particular, with probability at least $1/c$, p arrives at a random time step in the range $(t - n/c, t]$, in which case t appears to p to be a random step in the range $[1, n/c]$. Since the expected number of broadcasts by p in its first n/c steps is $\omega(\log n)$, the probability of p broadcasting at time t is $\omega\left(\frac{\log n}{n}\right)$.

For a given time step $t \in (n/c, n]$, we have shown that the number of players that attempt to broadcast in t is a sum of independent indicator random variables with mean at least $\omega(\log n)$ (i.e., asymptotically greater than $\log n$). It follows that, with high probability in n , there are multiple broadcast attempts at time t . Since this is true for every $t \in (n/c, n]$, we find that with high probability in n , the algorithm \mathcal{A} fails to achieve any successes in the first n steps. This contradicts the fact that \mathcal{A} is $f(n)$ -resilient to jamming. \square

REFERENCES

- [1] 2016. IEEE Standard for Information Technology–Telecommunications and Information Exchange Between Systems Local and Metropolitan Area Networks – Specific Requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. *IEEE Std 802.11-2016 (Revision of IEEE Std 802.11-2012)* (2016), 1–3534.
- [2] David J. Aldous. 1987. Ultimate Instability of Exponential Back-Off Protocol for Acknowledgment-Based Transmission Control of Random Access Communication Channels. *IEEE Trans. on Inform. Theory* IT-33, 2 (March 1987), 219–223.
- [3] Lakshmi Anantharamu, Bogdan S. Chlebus, and Mariusz A. Rokicki. 2009. Adversarial Multiple Access Channel with Individual Injection Rates. In *Proceedings of the 13th International Conference on Principles of Distributed Systems (OPODIS)*. 174–188.
- [4] William C. Anderton and Maxwell Young. 2017. Is Our Model for Contention Resolution Wrong?: Confronting the Cost of Collisions. In *Proceedings of the 29th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA*. 183–194.
- [5] Antonio Fernández Anta, Miguel A. Mosteiro, and Jorge Ramón Muñoz. 2013. Unbounded Contention Resolution in Multiple-Access Channels. *Algorithmica* 67, 3 (2013), 295–314.
- [6] Baruch Awerbuch, Andrea Richa, and Christian Scheideler. 2008. A Jamming-Resistant MAC Protocol for Single-Hop Wireless Networks. In *Proceedings of the 27th ACM Symposium on Principles of Distributed Computing (PODC)*. 45–54.

- [7] Michael A. Bender, Martin Farach-Colton, Simai He, Bradley C. Kuzmaul, and Charles E. Leiserson. 2005. Adversarial Contention Resolution for Simple Channels. In *Proceedings of the 17th Annual ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*. 325–332.
- [8] Michael A. Bender, Jeremy T. Fineman, and Seth Gilbert. 2006. Contention Resolution with Heterogeneous Job Sizes. In *Proceedings of the 14th Annual European Symposium on Algorithms (ESA)*. 112–123.
- [9] Michael A. Bender, Jeremy T. Fineman, Seth Gilbert, and Maxwell Young. 2016. How to Scale Exponential Backoff: Constant Throughput, Polylog Access Attempts, and Robustness. In *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.
- [10] Michael A. Bender, Jeremy T. Fineman, Seth Gilbert, and Maxwell Young. 2019. Scaling Exponential Backoff: Constant Throughput, Polylogarithmic Channel-Access Attempts, and Robustness. *J. ACM* 66, 1 (2019), 6:1–6:33.
- [11] Michael A. Bender, Tsvi Kopelowitz, William Kuzmaul, and Seth Pettie. 2020. Contention Resolution without Collision Detection. *ArXiv* (2020).
- [12] Michael A. Bender, Tsvi Kopelowitz, Seth Pettie, and Maxwell Young. 2016. Contention resolution with log-logstar channel accesses. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC*. 499–508.
- [13] D. J. Bernstein. 1998. qmail – An email Message Transfer Agent. <http://cr.ypt.to/qmail.html>.
- [14] Giuseppe Bianchi. 2006. Performance Analysis of the IEEE 802.11 Distributed Coordination Function. *IEEE Journal on Selected Areas in Communications* 18, 3 (Sept. 2006), 535–547.
- [15] Yi-Jun Chang, Wenyu Jin, and Seth Pettie. 2019. Simple Contention Resolution via Multiplicative Weight Updates. In *2nd Symposium on Simplicity in Algorithms (SOSA) (OASICS)*, Vol. 69. 16:1–16:16.
- [16] Yi-Jun Chang, Tsvi Kopelowitz, Seth Pettie, Ruosong Wang, and Wei Zhan. 2017. Exponential separations in the energy complexity of leader election. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC*. 771–783.
- [17] Bogdan S. Chlebus, Leszek Gasieniec, Dariusz R. Kowalski, and Tomasz Radzik. 2005. On the Wake-up Problem in Radio Networks. In *Proceedings of the 32nd International Colloquium on Automata, Languages and Programming (ICALP)*. 347–359.
- [18] Bogdan S. Chlebus and Dariusz R. Kowalski. 2004. A Better Wake-up in Radio Networks. In *Proceedings of 23rd ACM Symposium on Principles of Distributed Computing (PODC)*. 266–274.
- [19] Bogdan S. Chlebus, Dariusz R. Kowalski, and Mariusz A. Rokicki. 2006. Adversarial queuing on the multiple-access channel. In *Proc. Twenty-Fifth Annual ACM Symposium on Principles of Distributed Computing (PODC)*. 92–101.
- [20] Bogdan S. Chlebus, Dariusz R. Kowalski, and Mariusz A. Rokicki. 2012. Adversarial Queuing on the Multiple Access Channel. *ACM Transactions on Algorithms* 8, 1 (2012), 5.
- [21] Marek Chrobak, Leszek Gasieniec, and Dariusz R. Kowalski. 2007. The Wake-up Problem in Multihop Radio Networks. *SIAM J. Comput.* 36, 5 (2007), 1453–1471.
- [22] Bryan Costales and Eric Allman. 2002. *Sendmail* (third ed.). O'Reilly.
- [23] Gianluca De Marco, Dariusz R. Kowalski, and Grzegorz Stachowiak. 2018. Brief Announcement: Deterministic Contention Resolution on a Shared Channel. In *32nd International Symposium on Distributed Computing (DISC 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- [24] Gianluca De Marco and Grzegorz Stachowiak. 2017. Asynchronous shared channel. In *Proceedings of the ACM Symposium on Principles of Distributed Computing*. ACM, 391–400.
- [25] Pawel Garncarek, Tomasz Jurdzinski, and Dariusz R. Kowalski. 2018. Local Queuing Under Contention. In *32nd International Symposium on Distributed Computing (DISC 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- [26] Mihály Geréb-Graus and Thanasis Tsantilas. 1992. Efficient Optical Communication in Parallel Computers. In *Proceedings of the 4th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA)*. 41–48.
- [27] Leslie Ann Goldberg, Mark Jerrum, Tom Leighton, and Satish Rao. 1997. Doubly Logarithmic Communication Algorithms for Optical-Communication Parallel Computers. 26, 4 (Aug. 1997), 1100–1119.
- [28] Leslie Ann Goldberg and Philip D. MacKenzie. 1996. Analysis of Practical Backoff Protocols for Contention Resolution with Multiple Servers. In *Proc. Seventh Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 554–563.
- [29] Leslie Ann Goldberg, Philip D. Mackenzie, Mike Paterson, and Aravind Srinivasan. 2000. Contention Resolution with Constant Expected Delay. *J. ACM* 47, 6 (2000), 1048–1096.
- [30] Leslie Ann Goldberg, Yossi Matias, and Satish Rao. 1999. An Optical Simulation of Shared Memory. 28, 5 (Oct. 1999), 1829–1847.
- [31] Jonathan Goodman, Albert G. Greenberg, Neal Madras, and Peter March. 1988. Stability of Binary Exponential Backoff. *J. ACM* 35, 3 (July 1988), 579–602.
- [32] Google. 2014. GCM (Google Cloud Messaging) Advanced Topics. <http://developer.android.com/google/gcm/adv.html#retry>
- [33] Albert G. Greenberg, Philippe Flajolet, and Richard E. Ladner. 1987. Estimating the Multiplicities of Conflicts to Speed Their Resolution in Multiple Access Channels. *J. ACM* 34, 2 (April 1987), 289–325.
- [34] Albert G. Greenberg and Shmuel Winograd. 1985. A Lower Bound on the Time Needed in the Worst Case to Resolve Conflicts Deterministically in Multiple Access Channels. *JACM* 32, 3 (July 1985), 589–596.
- [35] Johan Håstad, Frank Thomson Leighton, and Brian Rogoff. 1996. Analysis of Backoff Protocols for Multiple Access Channels. *SIAM J. Comput.* 25, 4 (1996), 740–774. <https://doi.org/10.1137/S0097539792233828>
- [36] Maurice Herlihy and J. Eliot B. Moss. 1993. Transactional Memory: Architectural Support for Lock-Free Data Structures. In *Proceedings of the 20th International Conference on Computer Architecture*. 289–300. <http://www.cs.brown.edu/people/mph/isca2.ps>
- [37] V. Jacobson. 1988. Congestion Avoidance and Control. *SIGCOMM Comput. Commun. Rev.* 18, 4 (Aug. 1988), 314–329.
- [38] Colin McDiarmid. 1989. On the method of bounded differences. *Surveys in combinatorics* 141, 1 (1989), 148–188.
- [39] Robert M. Metcalfe and David R. Boggs. 1976. Ethernet: Distributed Packet Switching for Local Computer Networks. *Commun. ACM* 19, 7 (July 1976), 395–404.
- [40] Amit Mondal and Aleksandar Kuzmanovic. 2008. Removing Exponential Backoff from TCP. *SIGCOMM Comput. Commun. Rev.* 38, 5 (Sept. 2008), 17–28.
- [41] Adrian Ogierman, Andrea Richa, Christian Scheideler, Stefan Schmid, and Jin Zhang. 2018. Sade: competitive MAC under adversarial SINR. *Distributed Computing* 31, 3 (01 Jun 2018), 241–254.
- [42] Google Apps Platform. 2011. Google Documents List API version 3.0: Implementing Exponential Backoff. https://developers.google.com/google-apps/documents-list/?csw=1#implementing_exponential_backoff
- [43] Prabhakar Raghavan and Eli Upfal. 1999. Stochastic Contention Resolution With Short Delays. 28, 2 (April 1999), 709–719.
- [44] Ravi Rajwar and James R. Goodman. 2001. Speculative Lock Elision: Enabling Highly Concurrent Multithreaded Execution. In *Proc. of the 34th Annual Intl. Symposium on Microarchitecture*. Austin, Texas, 294–305. <http://www.cs.wisc.edu/~rajwar/papers/micro01.pdf>
- [45] Andrea Richa, Christian Scheideler, Stefan Schmid, and Jin Zhang. 2010. A Jamming-Resistant MAC Protocol for Multi-Hop Wireless Networks. In *Proceedings of the International Symposium on Distributed Computing (DISC)*. 179–193.
- [46] Andrea Richa, Christian Scheideler, Stefan Schmid, and Jin Zhang. 2011. Competitive and Fair Medium Access Despite Reactive Jamming. In *Proceedings of the 31st International Conference on Distributed Computing Systems (ICDCS)*. 507–516.
- [47] Andrea Richa, Christian Scheideler, Stefan Schmid, and Jin Zhang. 2012. Competitive and Fair Throughput for Co-Existing Networks Under Adversarial Interference. In *Proceedings of the 31st ACM Symposium on Principles of Distributed Computing (PODC)*. 291–300.
- [48] Andrea Richa, Christian Scheideler, Stefan Schmid, and Jin Zhang. 2013. Competitive Throughput in Multi-Hop Wireless Networks Despite Adaptive Jamming. *Distributed Computing* 26, 3 (2013), 159–171.
- [49] Andrea Richa, Christian Scheideler, Stefan Schmid, and Jin Zhang. 2013. An Efficient and Fair MAC Protocol Robust to Reactive Interference. *IEEE/ACM Transactions on Networking* 21, 1 (2013), 760–771.
- [50] Amazon Web Services. 2012. Error Retries and Exponential Backoff in AWS. <http://docs.aws.amazon.com/general/latest/gr/api-retries.html>
- [51] Nah-Oak Song, Byung-Jae Kwak, and Leonard E. Miller. 2003. On the Stability of Exponential Backoff. *Journal of Research of the National Institute of Standards and Technology* 108, 4 (2003).
- [52] Dan E. Willard. 1986. Log-logarithmic Selection Resolution Protocols in a Multiple Access Channel. *SIAM J. Comput.* 15, 2 (May 1986), 468–477.