An Improved Stability Bound for Binary Exponential Backoff

Hesham Al-Ammal* Leslie Ann Goldberg† Phil MacKenzie‡

February 24, 2000

Abstract

Goodman, Greenberg, Madras and March gave a lower bound of $n^{-\Omega(\log n)}$ for the maximum arrival rate for which the n-user binary exponential backoff protocol is stable. Thus, they showed that the protocol is stable as long as the arrival rate is at most $n^{-\Omega(\log n)}$. We improve the lower bound, showing that the protocol is stable for arrival rates up to $O(n^{-(.75+\delta)})$, for any $\delta > 0$.

1 Introduction

A *multiple-access channel* is a broadcast channel that allows multiple users to communicate with each other by sending messages onto the channel. If two or more users simultaneously send messages, then the messages interfere with each other (collide), and the messages are not transmitted successfully. The channel is not centrally controlled. Instead, the users use a contention-resolution protocol to resolve collisions. Thus, after a collision, each user involved in the collision waits a random amount of time (which is determined by the protocol) before re-sending. Perhaps the best-known contention-resolution protocol is the *Ethernet* protocol of Metcalfe and Boggs [10]. The Ethernet protocol is based on the following simple *binary exponential backoff protocol.* Time is divided into discrete units called steps. If the i'th user has a message to send during a given step, then it sends this message with probability 2^{-b_i}, where b_i denotes the number of collisions that this message has already had. With probability $1 - 2^{-b_i}$, user i does not send during the step. The Ethernet protocol is based on binary exponential backoff, but some modifications are made to make it easier to build. See [7, 10] for details.

Håstad, Leighton and Rogoff [7] have studied the performance of the binary exponential backoff protocol in the following natural model. The system consists of n users. Each user maintains a queue of messages that it wishes to send. At the beginning of the t'th time step, the length of the queue of the i'th user is denoted $q_i(t)$ and the number of times that the message at the head of its queue has collided is denoted $b_i(t)$. At the beginning of the t'th step, each queue receives 0 or 1 new messages. In particular, a new message is added

*hesham@dcs.warwick.ac.uk, Department of Computer Science, University of Warwick, Coventry, CV4 7AL, United Kingdom.

†leslie@dcs.warwick.ac.uk, http://www.dcs.warwick.ac.uk/~leslie/, Department of Computer Science, University of Warwick, Coventry, CV4 7AL, United Kingdom. This work was partially supported by EPSRC grant GR/L6098 and by ESPRIT Projects RAND-APX and ALCOM-FT.

‡philmac@research.bell-labs.com, Information Sciences Center, Bell Laboratories, Lucent Technologies, 600 Mountain Avenue, Murray Hill, NJ 07974–0636.

to the end of each queue independently with probability λ/n, where λ is the *arrival rate* of the system. (Thus, the length of the i'th queue is now $q_i(t) + Z$, where Z is a Bernoulli random variable, which is 1 with probability λ/n and 0 otherwise.) After the new messages are added to the queues, each user makes an independent decision about whether or not to send the message at the head of its queue, using the binary exponential backoff protocol. (If the message at the head of the i'th queue has never been sent before then $b_i = 0$, so it is now sent. Otherwise, $b_i = b_i(t)$, so it is sent independently with probability $2^{-b_i(t)}$.) If exactly one message is sent (so there are no collisions), then this message is delivered successfully, and it leaves its queue. Otherwise, the messages that are sent collide and no messages are delivered successfully.

Since the arrivals are modelled by a stochastic process, the evolution of the whole system over time can be viewed as a Markov chain in which the state just before step t is $X(t) = ((q_1(t), \ldots, q_n(t)), (b_1(t), \ldots, b_n(t)))$ and the next state is $X(t+1)$. The start state of the chain, $X(0)$, is $((0, \ldots, 0), (0, \ldots, 0))$. The chain is said to be *recurrent* if, with probability 1, it returns to its start state. That is, it is recurrent if

$$\Pr(X(t) = X(0) \text{ for some } t \geq 1) = 1.$$

It is said to be *positive recurrent* if the expected time that it takes to return to the start state is finite. In particular, let

$$T_{\text{ret}} = \min\{t \geq 1 \mid X(t) = X(0)\}.$$

The chain is said to be positive recurrent if $E[T_{\text{ret}}] < \infty$. Note that if the chain is not positive recurrent then the protocol is not a very good one. Informally, once it enters a "bad" state (one with a large backlog of messages), the expected time that it takes to get back to a state which is not bad is infinite. For this reason, we say that a protocol is *stable* if and only if the corresponding Markov chain is positive recurrent.[1] Håstad et al. [7] proved that if the arrival rate is too high, then the binary exponential backoff protocol is unstable, in the sense that the corresponding Markov chain is not positive recurrent.

Theorem 1 (Håstad, Leighton, and Rogoff) *Suppose that for some positive ϵ, $\lambda \geq \frac{1}{2}+\epsilon$. Suppose that n is sufficiently large (as a function of ϵ). Then $E[T_{ret}] = \infty$.*

On the other hand, Goodman, Greenberg, Madras and March [5] showed that if the arrival rate is sufficiently low, then the protocol is stable.

Theorem 2 (Goodman, Greenberg, Madras and March) *There is a positive constant α such that $E[T_{ret}]$ is finite for the n-user system, provided that $\lambda < \frac{1}{n^{\alpha \log n}}$.*

While Goodman, Greenberg, Madras, and March's result is the only known stability result for the finitely-many-users binary-exponential-backoff protocol, their upper bound ($\lambda < \frac{1}{n^{\alpha \log n}}$) is very small. In this paper, we narrow the gap between the two results. In particular, we prove the following theorem.

Theorem 3 *There is a positive constant α such that for any $\eta < 0.25$, as long as n is sufficiently large and $\lambda < \frac{1}{\alpha n^{1-\eta}}$ then $E[T_{ret}]$ is finite for the n-user system.*

[1]For further information about Markov chains, recurrence, positive recurrence and stability, see [2] and Chapter 6 of [6].

The main point of Theorem 3 is to show that n-user Binary Exponential Backoff is stable for an arrival rate that is the inverse of a polynomial in n, and in fact the inverse of a sublinear polynomial in n. With our specific proof technique, it seems that we cannot prove stability for rates higher than about $n^{-.75}$, and thus a natural open problem is to improve this bound. Perhaps the most interesting (and difficult) question raised by this work is whether an n-user system is stable for some constant arrival rate. For further discussion about improving our result, see Section 4.

The organisation of the paper is as follows. In Section 2 we summarise other related work. In Section 3 we give the proof of Theorem 3.

2 Related Work

We now summarize some other related work. We start by observing that the results in Theorem 1 and 2 can be extended to more general models. For example, the result of Goodman et al. can be extended to a more general model of stochastic arrivals in which the expected number of arrivals at user i at time t (conditioned on all events up to time t) is a quantity, λ_i, and $\sum_i \lambda_i$ is required to be equal to λ. The result of Håstad et al. can be extended to small values of n, provided that $\lambda > .568 + 1/(4n-2)$. The instability result of Håstad et al. implies that, when λ is sufficiently large, the expected average waiting time of messages is infinite.

Next, we mention that the binary exponential backoff protocol is known to be unstable in the infinitely-many-users Poisson-arrivals model. Kelly and MacPhee [8, 9] showed this for $\lambda > \ln 2$ and Aldous [1] showed that it holds for all positive λ.[2]

While the goal of this paper is to understand the binary-exponential backoff protocol, on which Ethernet is based, there are other acknowledgement-based protocols which are known to be stable in the same model for larger arrival rates. In particular, Håstad et al. have shown that *polynomial-backoff* protocols are stable as long as $\lambda < 1$. The expected waiting time of messages is high in polynomial-backoff protocols, but Raghavan and Upfal [11] have given a protocol that is stable for $\lambda < 1/10$, in which the expected waiting time of every message is $O(\log n)$, provided that the users are given a reasonably good estimate of $\log n$. Finally, Goldberg, MacKenzie, Paterson and Srinivasan [4] have given a protocol that is stable for $\lambda < 1/e$, in which the expected average message waiting-time is $O(1)$, provided that the users are given an upper bound on n.

We conclude by observing that the technique of Goldberg and MacKenzie [3] can be used to extend Theorem 3 so that it applies to a non-geometric version of binary-exponential backoff, which is closer to the version used in the Ethernet. (Instead of deciding whether to send on each step independently with probability 2^{-b_i}, the user simply chooses the number of steps to wait before sending uniformly at random from $[1, \ldots, 2^{b_i}]$.) The ideas are the same as those used in the proof that follows, but the details are messier. Our result can also be extended along the lines of [7] to show that, when λ is sufficiently low, the expected average message waiting time is finite.[3]

[2]Note that it can be misleading to view the infinitely-many-users model as the limit (as n tends to infinity) of the n-users model. For example, the "polynomial backoff" protocol is known to be unstable (for any positive λ) in the infinitely-many-users Poisson-arrivals model [8, 9], but it is stable (for any $\lambda < 1$) in the n-users model [7]. Thus, Aldous's result does not rule out the possibility that there is a positive constant λ^* such that the n-user binary exponential backoff protocol is stable whenever $\lambda < \lambda^*$.

[3]The word "stable" is not used consistently in the literature. For example, [7] incorporates the expected

3 The stability proof

In this section, we will prove Theorem 3. Let α be a sufficiently large positive constant and let η be a constant in the range $(0, .25)$. Suppose that the arrival rate λ is $\frac{1}{\alpha' n^{1-\eta}}$ for some $\alpha' \geq \alpha$. We will show that, if n is sufficiently large, the Markov chain corresponding to the binary exponential backoff protocol is positive recurrent.

The most common tool for proving that a Markov chain is positive recurrent is Foster's theorem.[4]

Theorem 4 (Foster) *A time-homogeneous irreducible aperiodic Markov chain X with a countable state space $\mathcal{A}$ is positive recurrent iff there exists a positive function $f(\rho)$, $\rho \in \mathcal{A}$, a number $\epsilon > 0$, and a finite set $C \subseteq \mathcal{A}$, such that the following inequalities hold.*

$$E[f(X(t+1)) - f(X(t)) \mid X(t) = \rho] \quad \leq \quad -\epsilon, \quad \rho \notin C \tag{1}$$

$$E[f(X(t+1)) \mid X(t) = \rho] \quad < \quad \infty, \quad \rho \in C. \tag{2}$$

Basically, the idea is to use a "potential function" f to follow the progress of the chain. The chain is positive recurrent iff there is a potential function f which

1. usually decreases (Equation 1), and
2. cannot increase much (Equation 2)

in a single step. Equation 1 implies that, from any state $\rho \notin C$, the expected time to reach C from ρ is at most $f(\rho)/\epsilon$. This (combined with Equation 2) implies that the expected return time to C is finite, which in turn implies that the chain is positive recurrent. (For more details, see [2].)

In practice, it can be difficult to find a potential function satisfying the criteria in Foster's theorem. We will use the following generalisation of the theorem due to Fayolle, Malyshev and Menshikov [2].

Theorem 5 (Fayolle, Malyshev, Menshikov) *A time-homogeneous irreducible aperiodic Markov chain X with a countable state space $\mathcal{A}$ is positive recurrent iff there exists a positive function $f(\rho)$, $\rho \in \mathcal{A}$, a number $\epsilon > 0$, a positive integer-valued function $k(\rho)$, $\rho \in \mathcal{A}$, and a finite set $C \subseteq \mathcal{A}$, such that the following inequalities hold.*

$$E[f(X(t+k(X(t)))) - f(X(t)) \mid X(t) = \rho] \quad \leq \quad -\epsilon k(\rho), \quad \rho \notin C \tag{3}$$

$$E[f(X(t+k(X(t)))) \mid X(t) = \rho)] \quad < \quad \infty, \quad \rho \in C. \tag{4}$$

---

average waiting time into the definition of "stability". Recall that in this paper, as in [5], stability means positive recurrence.

[4] A Markov chain is said to be *time-homogeneous* if its transition probabilities are fixed (for all time). It is *irreducible* if, for every pair of states (x, y), it is possible, in some number of steps, for the chain to move from state x to state y. It is *aperiodic* if, for any state x, the greatest common divisor of the set

$$\{t \mid \text{the chain can move from state } x \text{ to state } x \text{ in exactly } t \text{ steps}\}$$

is one. See [6] for details. The Markov chain corresponding to the binary exponential backoff protocol is time-homogeneous, irreducible, and aperiodic.

The reason that the generalisation is easier to use than Foster's theorem is that, while it may be difficult to find a potential function f which (usually) goes down in a single step, it may be easier to find one which goes down over several steps. In the generalised version of the theorem, it is only necessary to show that from a state ρ, the potential goes down by a factor of k over k steps, where k is allowed to depend upon ρ.

We will now define the potential function that we will use. The value μ in the potential function is a constant in the range $[\eta, 0.5-\eta)$. Let $f(X(t))$ be the following function of the state just before step t.

$$f(X(t)) = \alpha n^{2-\eta-\mu} \sum_{i=1}^{n} q_i(t) + \sum_{i=1}^{n} 2^{b_i(t)}.$$

We use the following notation, where $\beta = 3$. For a state $X(t)$, let $m(X(t))$ denote the number of users i with $q_i(t) > 0$ and $b_i(t) < \lg \beta + \lg n$, and let $m'(X(t))$ denote the number of users i with $q_i(t) > 0$ and $b_i(t) < (1-\eta-\mu)\lg n + 1$. We will take ϵ to be $1 - 2/\alpha$ and C to be the set consisting of the single state $((0,\ldots,0),(0,\ldots,0))$. We define $k(((0,\ldots,0),(0,\ldots,0))) = 1$, so Equation 4 is satisfied. For every state $\rho \notin C$, we will define $k(\rho)$ in such a way that Equation 3 is also satisfied. We give the details in three cases.

3.1 Case 1: $m'(X(t)) = 0$ and $m(X(t)) < n^{1-\eta-\mu}$.

For every state ρ such that $m'(\rho) = 0$ and $m(\rho) < n^{1-\eta-\mu}$ we define $k(\rho) = 1$. We wish to show that, if $\rho \neq ((0,\ldots,0),(0,\ldots,0))$ and $X(t) = \rho$, then $E[f(X(t+1)) - f(X(t))] \leq -\epsilon$. First, we give some intuition as to why the potential f is expected to drop in a single step. In this case (since $m'(X(t)) = 0$) all users which have messages to send have *large* backoff counters. Furthermore (since $m(X(t)) < n^{1-\eta-\mu}$) most backoff counters (all but at most $n^{1-\eta-\mu}$) are very large. This means that collisions are fairly unlikely. The expected drop in f mainly comes from the fact that if user i does send (which happens with probability 2^{-b_i}) and succeeds (which is fairly likely), then f drops by $2^{b_i} - 1$. The proof that f is expected to go down comes from a careful analysis of a single step and uses the same general approach as the one used in the proof of Lemma 5.7 of [7]. For convenience, we use m as shorthand for $m(X(t))$ and we use ℓ to denote the number of users i with $q_i(t) > 0$. Without loss of generality, we assume that these are users $1,\ldots,\ell$. We use p_i to denote the probability that user i sends on step t. (So $p_i = 2^{-b_i(t)}$ if $i \in [1,\ldots,\ell]$ and $p_i = \lambda/n$ otherwise.) We let T denote $\prod_{i=1}^{n}(1-p_i)$ and we let S denote $\sum_{i=1}^{n} \frac{p_i}{1-p_i}$. Note that the expected number of successes at step t is ST. Let $I_{a,i,t}$ be the 0/1 indicator random variable which is 1 iff there is an arrival at user i during step t and let $I_{s,i,t}$ be the 0/1 indicator random variable which is 1 iff user i succeeds in sending a message at step t. Then

$$E[f(X(t+1)) - f(X(t))] = \alpha n^{2-\eta-\mu} \sum_{i=1}^{n} (E[I_{a,i,t}] - E[I_{s,i,t}]) + \sum_{i=1}^{n} \left(E[2^{b_i(t+1)}] - 2^{b_i(t)}\right),$$

$$= \alpha n^{2-\eta-\mu}\lambda - \alpha n^{2-\eta-\mu} ST + \sum_{i=1}^{n} \left(2^{b_i(t)}\sigma_i - (2^{b_i(t)} - 1)\pi_i\right), \quad (5)$$

$$= \alpha n^{2-\eta-\mu}\lambda - \alpha n^{2-\eta-\mu} ST + \sum_{i=1}^{n} \left(2^{b_i(t)} p_i (1 - \frac{T}{1-p_i}) - (2^{b_i(t)} - 1) p_i \frac{T}{1-p_i}\right),$$

$$= \alpha n^{2-\eta-\mu}\lambda - \alpha n^{2-\eta-\mu} ST + \sum_{i=1}^{\ell} (1 - \frac{T}{1-p_i}) + \sum_{i=\ell+1}^{n} \frac{\lambda}{n}(1 - \frac{T}{1-p_i}) - \ell T,$$

$$= \alpha n^{2-\eta-\mu}\lambda - \alpha n^{2-\eta-\mu} ST + \ell - \ell T + \frac{(n-\ell)\lambda}{n} - T\left(\sum_{i=1}^{\ell}\frac{1}{1-p_i} + \sum_{i=\ell+1}^{n}\frac{p_i}{1-p_i}\right),$$
$$= \alpha n^{2-\eta-\mu}\lambda - \alpha n^{2-\eta-\mu} ST + \ell - \ell T + \frac{(n-\ell)\lambda}{n} - ST - \ell T,$$
$$= \alpha n^{2-\eta-\mu}\lambda + \ell + \frac{(n-\ell)\lambda}{n} - T((\alpha n^{2-\eta-\mu}+1)S + 2\ell), \tag{6}$$

where σ_i in Equality 5 denotes the probability that user i collides at step t and π_i denotes the probability that user i sends successfully at step t. (To see why Equality 5 holds, note that with probability σ_i, $b_i(t+1) = b_i(t)+1$, with probability π_i, $b_i(t+1) = 0$, and otherwise, $b_i(t+1) = b_i(t)$.) We now find lower bounds for S and T. First,

$$\begin{aligned} S &= \sum_{i=1}^{n}\frac{p_i}{1-p_i} \\ &= \sum_{i=1}^{\ell}\left(\frac{2^{-b_i(t)}}{1-2^{-b_i(t)}}\right) + \frac{\lambda(n-\ell)}{n-\lambda} \\ &\geq \sum_{i=1}^{m}\left(\frac{1}{\beta n - 1}\right) + \frac{\lambda(n-\ell)}{n-\lambda} \\ &= \frac{m}{\beta n - 1} + \frac{\lambda(n-\ell)}{n-\lambda}. \end{aligned} \tag{7}$$

Next,

$$\begin{aligned} T &= \prod_{i=1}^{n}(1-p_i) \\ &\geq (1 - \frac{1}{2n^{1-\eta-\mu}})^{m}(1-\frac{1}{\beta n})^{\ell-m}(1-\frac{\lambda}{n})^{n-\ell} \\ &\geq 1 - \frac{m}{2n^{1-\eta-\mu}} - \frac{\ell-m}{\beta n} - \frac{\lambda(n-\ell)}{n} \end{aligned} \tag{8}$$

Combining Equations 6, 7 and 8, we get the following equation.

$$E[f(X(t+1)) - f(X(t))] \leq \alpha n^{2-\eta-\mu}\lambda + \ell + \frac{(n-\ell)\lambda}{n} - \tag{9}$$
$$\left(1 - \frac{m}{2n^{1-\eta-\mu}} - \frac{\ell-m}{\beta n} - \frac{\lambda(n-\ell)}{n}\right)\left((\alpha n^{2-\eta-\mu}+1)\left(\frac{m}{\beta n - 1} + \frac{\lambda(n-\ell)}{n-\lambda}\right) + 2\ell\right).$$

We will let $g(m,\ell)$ be the quantity in Equation 9 plus ϵ and we will show that $g(m,\ell)$ is negative for all values of $0 \leq m < n^{1-\eta-\mu}$ and all $\ell \geq m$. In particular, for every fixed positive value of m, we will show that

1. $g(m,m)$ is negative,
2. $g(m,n)$ is negative, and
3. $\frac{\partial^2}{\partial \ell^2} g(m,\ell) > 0$. ($g(m,\ell)$ is concave up as a function of ℓ for the fixed value of m so $g(m,\ell)$ is negative for all $\ell \in [m,n]$.)

We will handle the case $m = 0$ similarly except that $m = \ell = 0$ corresponds to the start state, so we will replace Item 1 with the following for $m = 0$.

1'. $g(0, 1)$ is negative.

The details of the proof are now merely calculations.

1. $g(m, m)$ **is negative:**

$$
\begin{aligned}
&g(m,m) \times 2\,\alpha'\,n^{(2-2\,\eta)}\,(\beta\,n-1)\,(\alpha'\,n^2-n^\eta) = \\
&\quad n^{(-2\,\eta+5)}\,m^2\,\alpha\,\alpha'^2 + 2\,n - 2\,m - 2\,n^{(-2\,\eta+4)}\,\varepsilon\,\alpha'^2 - 2\,n^{(-2\,\eta+5)}\,\alpha'^2\,m\,\beta \\
&\quad + 2\,\beta\,n^3 + 2\,n^{(-2\,\eta+5)}\,\varepsilon\,\beta\,\alpha'^2 - 2\,n^{(-\eta+3)}\,\varepsilon\,\alpha'\,\beta - n^{(-\eta+4)}\,m^2\,\alpha\,\beta\,\alpha' \\
&\quad - 2\,n^{(-\eta-\mu+4)}\,\alpha\,\beta - 4\,m^2 - 2\,n^{(-3\,\eta-\mu+6)}\,\alpha\,m\,\alpha'^2 + 2\,n^{(-2\,\eta-\mu+5)}\,m\,\alpha\,\alpha' + 6\,m\,n \\
&\quad + 2\,n^{(-2\,\eta-\mu+5)}\,m\,\alpha\,\alpha'\,\beta - 2\,n^{(-2\,\eta-\mu+4)}\,m^2\,\alpha\,\alpha' + 2\,n^{(-\eta+3)}\,m\,\beta\,\alpha' \\
&\quad - n^{(-\eta+4)}\,m\,\alpha\,\alpha' - 2\,n^2 + 2\,n^{(-\eta-\mu+5)}\,\alpha\,\beta - 2\,\beta\,n^2 + 2\,n^{(-\eta-\mu+3)}\,\alpha \\
&\quad + n^{(\mu+3)}\,m\,\alpha'\,\beta + 6\,m^2\,\beta\,n - 8\,\beta\,n^2\,m + n^{(-\eta+5)}\,m\,\alpha\,\beta\,\alpha' + 2\,n^{(-\eta-\mu+3)}\,\alpha\,m \\
&\quad - 2\,n^{(-\eta+3)}\,m\,\alpha' - 3\,n^{(\mu+2)}\,m^2\,\alpha'\,\beta + 4\,n^{(-\eta+4)}\,m\,\beta\,\alpha' + 2\,n^{(-\eta-\mu+3)}\,m^2\,\alpha\,\beta \\
&\quad - 2\,n^{(-\eta-\mu+4)}\,\alpha + 2\,n^{(-\eta+2)}\,\varepsilon\,\alpha' + 2\,m^2\,n^{(\mu+1)}\,\alpha' + 2\,n^{(-\eta+\mu+4)}\,m^2\,\beta\,\alpha'^2 \\
&\quad - 4\,n^{(-\eta-\mu+4)}\,\alpha\,m\,\beta - 4\,n^{(-\eta+3)}\,m^2\,\beta\,\alpha' - n^{(-\eta+\mu+3)}\,m^2\,\alpha'^2 + 2\,n^{(-\eta+2)}\,m^2\,\alpha' \\
&\quad - n^{(\mu+2)}\,m\,\alpha' + 2\,\beta\,n\,m
\end{aligned}
$$

Since $\eta + \mu < .5$, the dominant term is $-2n^{(-3\eta-\mu+6)}\alpha m\alpha'^2$. Note that there is a positive term ($n^{(-2\eta+5)}m^2\alpha\alpha'^2$) which could be half this big if m is as big as $n^{1-\eta-\mu}$ (the upper bound for Case 1), but all other terms are asymptotically smaller.

2. $g(m, n)$ **is negative:**

$$
\begin{aligned}
&g(m,n) \times 2\,\alpha'\,\beta\,n\,(\beta\,n-1) = \\
&\quad 2\,\alpha\,n^{(-\mu+3)}\,\beta^2 - 2\,\alpha\,n^{(-\mu+2)}\,\beta - 2\,\alpha'\,\beta^2\,n^3 + 2\,\alpha'\,\beta\,n^2 \\
&\quad - 2\,m\,\alpha\,n^{(-\eta-\mu+3)}\,\alpha'\,\beta - 2\,\alpha'\,\beta\,n\,m + \alpha'\,\beta\,n^2\,m^2\,\alpha + m^2\,n^{(\eta+\mu)}\,\alpha'\,\beta \\
&\quad + 2\,m\,n^{(\eta+\mu+2)}\,\alpha'\,\beta^2 - 2\,m\,n^{(\eta+\mu+1)}\,\alpha'\,\beta + 2\,m\,\alpha\,n^{(-\eta-\mu+3)}\,\alpha' + 6\,\alpha'\,n\,m \\
&\quad + 4\,\alpha'\,\beta\,n^3 - 4\,\alpha'\,n^2 - 2\,n^{(-\eta-\mu+2)}\,m^2\,\alpha\,\alpha' - 2\,\alpha'\,m^2 - 4\,\alpha'\,\beta\,n^2\,m + 2\,\varepsilon\,\alpha'\,\beta^2\,n^2 \\
&\quad - 2\,\varepsilon\,\alpha'\,\beta\,n
\end{aligned}
$$

Since $\eta + \mu < .5$ and $\beta > 2$, the term $-2\alpha'\beta^2n^3$ dominates $+4\alpha'\beta n^3$. For the same reason, the term $-2m\alpha n^{(-\eta-\mu+3)}\alpha'\beta$ dominates the two terms $+2m\alpha n^{(-\eta-\mu+3)}\alpha'$ and $+\alpha'\beta n^2m^2\alpha$. The other terms are asymptotically smaller.

3. $\frac{\partial^2}{\partial \ell^2} g(m, \ell) > 0$**:**

$$
\frac{\partial^2}{\partial \ell^2} g(m,\ell) = -2\,(-\frac{1}{\beta\,n} + \frac{\lambda}{n})\,(-\frac{(\alpha\,n^{(-\eta-\mu+2)}+1)\,\lambda}{n-\lambda} + 2)
$$

1'. $g(0,1)$ **is negative:**

$$
\begin{aligned}
& g(0,1) \times \alpha' \, \beta \, n^{(2-2\,\eta)} \, (\alpha' \, n^2 - n^\eta) = \\
& \quad -n^{(-2\,\eta-\mu+3)} \, \alpha \, \alpha' - n^{(-\eta+2)} \, \alpha' \, \beta + n^{(-\eta-\mu+2)} \, \alpha \, \beta + \beta \, n^2 \\
& \quad - n^{(-\eta+2)} \, \varepsilon \, \beta \, \alpha' + 2 \, n^{(-2\,\eta+3)} \, \alpha'^2 - n^{(-2\,\eta+4)} \, \alpha'^2 \, \beta + 2 \, n^{(-\eta+3)} \, \beta \, \alpha' \\
& \quad + n^{(-2\,\eta+4)} \, \varepsilon \, \beta \, \alpha'^2 + n^{(-\eta-\mu+4)} \, \alpha \, \beta + n^{(-\eta+2)} \, \alpha' + n^{(-2\,\eta-\mu+4)} \, \alpha \, \alpha' \\
& \quad + n^{(-2\,\eta-\mu+4)} \, \alpha \, \beta \, \alpha' + 4 \, \beta - 3 \, n^{(-\eta+1)} \, \alpha' - 5 \, \beta \, n - 3 \, n^{(-\eta-\mu+3)} \, \alpha \, \beta
\end{aligned}
$$

Since $\alpha'(1-\epsilon) \geq \alpha(1-\epsilon) > 1$, $\eta + \mu < .5$, and $\mu \geq \eta$, the term $-n^{(-2\eta+4)}\alpha'^2\beta(1-\epsilon)$ dominates the term $+n^{(-\eta-\mu+4)}\alpha\beta$. The other terms are asymptotically smaller.

3.2 Case 2: $m(X(t)) \geq n^{1-\eta-\mu}$ or $m'(X(t)) > n^{.4}$.

For every state ρ such that $m(\rho) \geq n^{1-\eta-\mu}$ or $m'(\rho) > n^{.4}$, we will define an integer k (which depends upon ρ) and we will show that, if $X(t) = \rho$, then $E[f(X(t+k)) - f(X(t))] \leq -\epsilon k$, where $\epsilon = 1 - 2/\alpha$.

For convenience, we will use m as shorthand for $m(X(t))$ and m' as shorthand for $m'(X(t))$. If $m \geq n^{1-\eta-\mu}$ then we will define $r = m$, $W = n^{\eta+\mu}\lceil \lg r \rceil 2^{-8}$, $A = W$, $b = \lg \beta + \lg n$ and $v = n$. Otherwise, we will define $r = m'$, $W = \lceil \lg r \rceil 2^{-8}$, $A = 0$, $b = (1-\eta-\mu)\lg n + 1$, and $v = 2\lceil n^{1-\eta-\mu} \rceil$. In either case, we will define $k = 4(r+v)\lceil \lg r \rceil$.

The intuition behind the proof is as follows. First, since many users have small backoff counters, it is fairly likely that a collision occurs on the first step. So we do not expect the potential f to drop in a single step. Instead, we study the evolution of the system over k steps. With sufficiently high probability, the backoff counters get driven up during the first $\Theta(r \log r)$ steps. (We refer to these steps as "the preamble".) During the remaining steps, the backoff counters stay reasonably high except during steps which occur shortly after

1. arrivals (but there are likely to be few of these since we only run for k steps), and
2. successful sends (which help to reduce f).

We refer to these as "exceptional steps". Without loss of generality, there are few of them, since otherwise there are many successes and the potential goes down. Although the backoff counters stay high (as we just explained), most of them do not get too high, since we only run for k steps. So the probability of success during any given step which is not exceptional or in the preamble is high. Finally, with sufficiently high probability, there are at least W successes, and this reduces the potential.

A technical difficulty in the proof is clarifying the independence between some of the events and for this reason, it is helpful to identify "preamble steps" (steps in τ_0), "exceptional steps" (steps in τ_1), and also "following steps". (The formal definition of "following steps" is given later. Typically, these steps follow at least W successes). The details of this partition of steps will be described later.

Let τ be the set of all steps $\{t, \ldots, t+k-1\}$ and let $\mathcal{S}$ be the random variable which denotes the number of successes that the system has during τ. Let p denote $\Pr(\mathcal{S} \geq W)$. Then we have

$$
E[f(X(t+k)) - f(X(t))] \quad \leq \quad \alpha n^{2-\eta-\mu}\lambda k - \alpha n^{2-\eta-\mu} E[\mathcal{S}] + \sum_{i=1}^{n} \sum_{t'=t+1}^{t+k} E[2^{b_i(t')} - 2^{b_i(t'-1)}]
$$

$$\leq \alpha n^{2-\eta-\mu}\lambda k - \alpha n^{2-\eta-\mu} W p + kn$$
$$\leq -\epsilon k,$$

where the final inequality holds as long as $\alpha p \geq 2^{13}$ and n is sufficiently big (see the Appendix). Thus, it suffices to find a positive lower bound for p which is independent of n. We do this with plenty to spare. In particular, we show that $p \geq 1 - 5 \times 10^{-5}$.

We start with a technical lemma, which describes the behaviour of a single user.

Lemma 6 *Let j be a positive integer, and let δ be a positive integer which is at least 2. Suppose that $q_i(t) > 0$. Then, with probability at least $1 - \frac{\lceil \lg j \rceil}{j^{\delta/(2\ln 2)}}$, either user i succeeds in at least one of the steps in the interval $[t, \ldots, t + \delta j \lceil \lg j \rceil - 1]$, or $b_i(t + \delta j \lceil \lg j \rceil) \geq \lceil \lg j \rceil$.*

Proof: Suppose that user i is running in an externally-jammed channel (so every send results in a collision). Let X_z denote the number of steps $t' \in [t, \ldots, t + \lceil \delta j \lg(j) \rceil]$ with $b_i(t') = z$. We claim that $\Pr(X_z > \delta \lceil \lg j \rceil 2^{z-1}) < j^{-\delta/(2\ln 2)}$. This proves the lemma since $\sum_{z=0}^{\lceil \lg j \rceil - 1} \delta \lceil \lg j \rceil 2^{z-1} \leq \delta j \lceil \lg j \rceil$. To prove the claim, note that $X_0 \leq 1$, so $\Pr(X_0 > \delta \lceil \lg j \rceil 2^{-1}) = 0 < j^{-\delta/(2\ln 2)}$. For $z > 0$, note that

$$\Pr(X_z > \delta \lceil \lg j \rceil 2^{z-1}) \leq (1 - 2^{-z})^{\delta \lceil \lg j \rceil 2^{z-1}} < j^{-\delta/(2\ln 2)}.$$

□

Next, we define some events. We will show that the events are likely to occur, and, if they do occur, then $\mathcal{S}$ is likely to be at least W. This will allow us to conclude that $p \geq 1 - 5 \times 10^{-5}$, which will finish Case 2. We start by defining $B = \lceil W \rceil + \lceil A \rceil$, $k' = 4r\lceil \lg r \rceil$, and $k'' = 4B\lceil \lg B \rceil$. Next, we give names to some of the steps in $\tau = \{t, \ldots, t + k - 1\}$. Let τ_0 be the preamble of τ consisting of steps $\{t, \ldots, t + k' - 1\}$. For every i, let $\tau'(i)$ be the set of times in τ when user i will "definitely" send. In particular, $t' \in \tau'(i)$ if and only if

1. $b_i(t') = 0$ and $q_i(t') > 0$, or
2. $b_i(t') = 0$ and there is an arrival at user i at t'.

τ_2 will be the suffix of following steps in τ. In particular, $t' \in \tau_2$ if and only if there are at least B pairs (t'', i) with $t'' < t'$ and $t'' \in \tau'(i)$. (Informally, by the time τ_2 is entered, there will have been at least B "definite sends", some of which may have coincided in time.) Note that τ_2 is a random variable. Finally, τ_1 will be a (possibly non-contiguous) subset of $\tau - \tau_0 - \tau_2$. Informally, τ_1 will contain all steps which occur during or shortly after "definite sends." Formally, τ_1 will be the set of all $t' \in \tau - \tau_0 - \tau_2$ such that, for some i, $\tau'(i) \cap [t' - k'' + 1, t'] \neq \emptyset$. See Figure 1.

We can now define the events E1–E4.

E1. There are at most A arrivals during τ.

E2. Every station with $q_i(t) > 0$ and $b_i(t) < b$ either sends successfully during τ_0 or has $b_i(t + k') \geq \lceil \lg r \rceil$.

E3. At least half of the stations with $q_i(t) > 0$ and $b_i(t) < b$ have $b_i(t') \leq b + \lceil \lg \lg(r) \rceil + 6$ for all $t' \in \tau$.

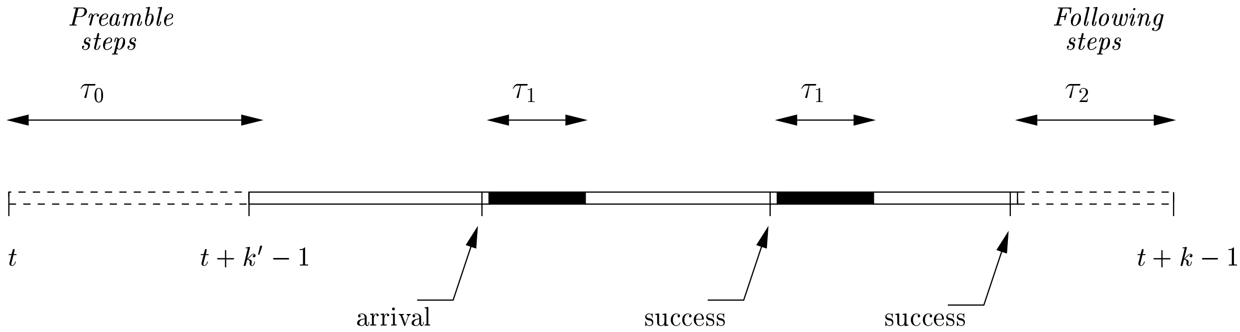

Figure 1: A possible outcome for the random variables τ_0, τ_1 and τ_2. For illustration, we assume that there are no arrivals or successes during the last k'' steps of the preamble.

E4. For all $t' \in \tau'(i)$ and all $t'' > t'$ such that $t'' \in \tau - \tau_0 - \tau_1 - \tau_2$, either $q_i(t'') = 0$ or $b_i(t'') \geq \lceil \lg B \rceil$.

Next, we show that E1–E4 are likely to occur.

Lemma 7 *If n is sufficiently large then* $\Pr(\overline{E1}) \leq 10^{-5}$.

Proof: The expected number of arrivals in τ is λk. If $m \geq n^{1-\eta-\mu}$, then $A = n^{\eta+\mu}\lceil \lg r \rceil 2^{-8} \geq 2\lambda k$. By a Chernoff bound, the probability that there are this many arrivals is at most $e^{-\lambda k/3} \leq 10^{-5}$. Otherwise, $A = 0$ and $\lambda k = o(1)$. Thus, $\Pr(\mathrm{E1}) \geq (1-\lambda/n)^{nk} \geq 1 - \lambda k \geq 1 - 10^{-5}$. □

Lemma 8 *If n is sufficiently large then* $\Pr(\overline{E2}) \leq 10^{-5}$.

Proof: Apply Lemma 6 to each of the r users with $\delta = 4$ and $j = r$. Then $\Pr(\overline{\mathrm{E2}}) \leq r\frac{\lceil \lg r \rceil}{r^{2/(\ln 2)}} \leq 10^{-5}$. □

Lemma 9 *If n is sufficiently large then* $\Pr(\overline{E3}) \leq 10^{-5}$.

Proof: Note that $k \leq 16v \lg r$. Also note that the probability of a given user i sending at step t' when $b_i(t') = b + \lceil \lg \lg(r) \rceil + 6$ is at most $1/(64v \lg r)$. Thus the probability that user i sends at all in the k steps of τ is at most $1/4$. By a Chernoff bound, the probability that over half of the r users with $q_i(t) > 0$ and $b_i(t) < b$ send when $b_i(t') = b + \lceil \lg \lg(r) \rceil + 6$ for some $t' \in \tau$ is at most $e^{-\Theta(r)} < 10^{-5}$. □

Lemma 10 *If n is sufficiently large then* $\Pr(\overline{E4}) \leq 10^{-5}$.

Proof: We can apply Lemma 6 separately to each of the (up to B) pairs (t', i) with $\delta = 4$ and $j = B$. The probability that event E4 does not hold is at most $\frac{B\lceil \lg B \rceil}{B^{2/(\ln 2)}} \leq 10^{-5}$. □

We now wish to show that $\Pr(\mathcal{S} < W \mid \mathrm{E1} \wedge \mathrm{E2} \wedge \mathrm{E3} \wedge \mathrm{E4}) \leq 10^{-5}$. We begin with the following lemma.

Lemma 11 *Given any fixed sequence of states $X(t), \ldots, X(t+z)$ which does not violate E2 or E4, and satisfies $t+z \in \tau - \tau_0 - \tau_1 - \tau_2$, $q_i(t+z) > 0$, and $b_i(t+z) \leq b + \lceil \lg\lg(r) \rceil + 6$, the probability that user i succeeds at step $t+z$ is at least $\frac{1}{2^{14}2^b \lg r}$.*

Proof: The conditions in the lemma imply the following.

- There are no users j with $q_j(t+z) > 0$ and $b_j(t+z) < \lceil \lg B \rceil$ (since E4 holds).
- There are at most B users j with $b_j(t+z) < \lceil \lg r \rceil$ (since E2 holds and at most B users succeed or have new arrivals).
- There are at most $r+B$ users j with $b_j(t+z) < b$ (since r started that way and at most B succeed or have new arrivals).
- There are at most $m+B$ users j with $b_j(t+z) < \lg \beta + \lg n$ (for similar reasons).

Thus, the probability that user i succeeds is at least

$$\begin{aligned}
& 2^{-(b+\lceil \lg\lg(r) \rceil+6)} \left(1 - \frac{1}{B}\right)^B \left(1 - \frac{1}{r}\right)^r \left(1 - \frac{1}{2^b}\right)^{m-r} \left(1 - \frac{1}{\beta n}\right)^{n-m-B} \\
\geq\ & \frac{1}{2^b \lg r 2^7} \frac{1}{4}\frac{1}{4}\frac{1}{4} \left(1 - \frac{n-m-B}{\beta n}\right) \\
\geq\ & \frac{1}{2^{14}2^b \lg r}.
\end{aligned}$$

□

Corollary 12 *Given any fixed sequence of states $X(t), \ldots, X(t+z)$ which does not violate E2, E3, or E4, and satisfies $t+z \in \tau - \tau_0 - \tau_1 - \tau_2$, the probability that some user succeeds at step $t+z$ is at least $\frac{(r/2)-B}{2^{14}2^b \lg r} \geq \frac{r}{2^{18} v \lg r}$.*

Proof: Since $t+z \notin \tau_2$, at least $r-B$ of the users i with $q_i(t) > 0$ and $b_i(t) < b$ have not succeeded before step $t+z$. Since E3 holds, at least $r/2 - B$ of these have $b_i(t+z) \leq b + \lceil \lg\lg(r) \rceil + 6$. For all i and i', the event that user i succeeds at step $t+z$ is disjoint with the event that user i' succeeds at step $t+z$. Finally, note that $(r/2) - B > r/4$ and $2^b \leq 4v$.
□

Lemma 13 *If n is sufficiently large then* $\Pr(\mathcal{S} < W \mid E1 \wedge E2 \wedge E3 \wedge E4) \leq 10^{-5}$.

Proof: If E1 is satisfied then τ_2 does not start until there have been at least W successes. Since $|\tau - \tau_0 - \tau_1| \geq k - k' - Bk'' \geq v\lceil \lg r \rceil/2$, Corollary 12 shows that the probability of having fewer than W successes is at most the probability of having fewer than W successes in $v\lceil \lg r \rceil/2$ Bernoulli trials with success probability $\frac{r}{2^{18} v \lg r}$. Since W is at most half of the expected number of successes, a Chernoff bound shows that the probability of having fewer than W successes is at most $\exp(-\frac{rv\lceil \lg r \rceil}{2^{22} v \lg r}) \leq 10^{-5}$. □

We conclude Case 2 by observing that p is at least $1 - \Pr(\overline{\mathrm{E1}}) - \Pr(\overline{\mathrm{E2}}) - \Pr(\overline{\mathrm{E3}}) - \Pr(\overline{\mathrm{E4}}) - \Pr(\mathcal{S} < W \mid \mathrm{E1} \wedge \mathrm{E2} \wedge \mathrm{E3} \wedge \mathrm{E4})$. By Lemmas 7, 8, 9, 10, and 13, this is at least $1 - 5 \times 10^{-5}$.

3.3 Case 3: $0 < m'(X(t)) \leq n^{.4}$ and $m(X(t)) < n^{1-\eta-\mu}$.

For every state ρ such that $0 < m'(\rho) \leq n^{.4}$ and $m(\rho) < n^{1-\eta-\mu}$, we will define $k = 32m'(\rho)\lceil \lg m'(\rho)\rceil + \lceil n^{1-\eta-\mu}\rceil$. We will show that, if $X(t) = \rho$, then $E[f(X(t+k)) - f(X(t))] \leq -\epsilon k$.

The intuition behind the proof in this case is similar to that of Case 2 except that we do not have enough small backoff counters to achieve W successes (as in Case 2) even though we may have too many to make the potential drop in a single step (as in Case 1). We study the evolution of the system over k steps. The backoff counters are likely to be driven up in the first $\Theta(m' \log m')$ steps. After that, we are likely to have a single success, which is enough to make the potential drop.

Once again, we will use m as shorthand for $m(X(t))$ and m' as shorthand for $m'(X(t))$. Let $\tau = \{t, \ldots, t+k-1\}$, let $\mathcal{S}$ be the number of successes that the system has in τ. Let p denote $\Pr(\mathcal{S} \geq 1)$. As in Case 2, $E[f(X(t+k)) - f(X(t))] \leq \alpha n^{2-\eta-\mu}\lambda k - \alpha n^{2-\eta-\mu} p + kn$, and this is at most $-\epsilon k$ as long as $\alpha p > 9$. Thus, we will finish by finding a positive lower bound for p which is independent of n.

Since $m' > 0$, there is a user γ such that $b_\gamma(t) < (1-\eta-\mu)\lg n + 1$. Let $k' = 32m'\lceil \lg m'\rceil$ and $\tau_0 = \{t, \ldots, t+k'-1\}$. We will now define some events, as in Case 2.

E1. There are no arrivals during τ.

E2. Every station with $q_i(t) > 0$ and $b_i(t) < (1-\eta-\mu)\lg n + 1$ either sends successfully during τ_0 or has $b_i(t+k') \geq \lceil \lg m'\rceil$.

E3. $b_\gamma(t') < (1-\eta-\mu)\lg n + 7$ for all $t' \in \tau$.

Lemma 14 *If n is sufficiently large then $Pr(\overline{E1}) \leq 10^{-5}$.*

Proof: As in the proof of Lemma 7,

$$\Pr(\mathrm{E1}) \geq \left(1 - \frac{\lambda}{n}\right)^{nk} \geq 1 - \lambda k \geq 1 - 10^{-5}.$$

□

Lemma 15 *$Pr(\overline{E2}) \leq 10^{-5}$.*

Proof: We use lemma 5 with $\delta = 32$ and $j = m'$ to get

$$Pr(\overline{E2}) \leq m' \cdot \frac{\lceil \lg m'\rceil}{(m')^{16/\ln(2)}} \leq 10^{-5}.$$

□

Lemma 16 *If n is sufficiently large then $Pr(\overline{E3}) \leq 10^{-5}$.*

Proof: For E3 to be violated, user γ must make at least 6 attempts, one each with backoff counter $\lceil (1-\eta-\mu)\lg n + r\rceil$ for $r \in \{1,\ldots,6\}$. The probability of this happening is

$$\begin{aligned} Pr(\overline{E3}) &\le \binom{k}{6}\prod_{r=1}^{6} 2^{-\lceil(1-\eta-\mu)\lg n\rceil - r} \\ &\le \left(\frac{ke}{6}\right)^6 \left(\frac{1}{n^{1-\eta-\mu}}\right)^6 2^{-\sum_{r=1}^{6} r} \\ &\le \left(\frac{2en^{1-\eta-\mu}}{6n^{1-\eta-\mu}2^3}\right)^6 \\ &\le 10^{-5}. \end{aligned}$$

□

Lemma 17 *Given any fixed sequence of states $X(t),\ldots,X(t+z)$ which does not violate E1, E2, or E3 such that $t+z \in \tau - \tau_0$ and there are no successes during steps $[t,\ldots,t+z-1]$, the probability that user γ succeeds at step $t+z$ is at least $\frac{1}{2^{12}n^{1-\eta-\mu}}$.*

Proof: The conditions in the statement of the lemma imply the following.

- $q_\gamma(t+z) > 0$ and $b_\gamma(t+z) < (1-\eta-\mu)\lg n + 7$.
- There are no users j with $b_j(t+z) < \lceil \lg m' \rceil$.
- There are at most m' users j with $b_j(t+z) < (1-\eta-\mu)\lg n + 1$.
- There are at most m users j with $b_j(t+z) < \lg\beta + \lg n$.
- There will be no arrivals on step $t+z$.

The probability of success for user γ is at least

$$\begin{aligned} & 2^{-((1-\eta-\mu)\lg n + 7)}\left(1-\frac{1}{m'}\right)^{m'-1}\left(1-\frac{1}{2n^{1-\eta-\mu}}\right)^{m-m'}\left(1-\frac{1}{\beta n}\right)^{n-m} \\ \ge\ & \frac{1}{2^7 n^{1-\eta-\mu}}\frac{1}{4}\frac{1}{4}\frac{1}{2} \\ \ge\ & \frac{1}{2^{12}n^{1-\eta-\mu}}. \end{aligned}$$

□

Lemma 18 *If n is sufficiently large then $Pr(\mathcal{S} < 1 \mid E1 \wedge E2 \wedge E3) \le e^{-1/2^{12}}$.*

Proof: Lemma 17 implies that the probability of having no successes is at most the probability of having no successes in $|\tau - \tau_0|$ Bernoulli trials, each with success probability $\frac{1}{2^{12}n^{1-\eta-\mu}}$. Since $|\tau - \tau_0| \ge n^{1-\eta-\mu}$, this probability is at most

$$\left(1 - \frac{1}{2^{12}n^{1-\eta-\mu}}\right)^{n^{1-\eta-\mu}} \le e^{-1/2^{12}}.$$

□

We conclude Case 3 by observing that

$$p \ge 1 - \Pr(\overline{\mathrm{E1}}) - \Pr(\overline{\mathrm{E2}}) - \Pr(\overline{\mathrm{E3}}) - \Pr(\mathcal{S} < 1 \mid \mathrm{E1} \wedge \mathrm{E2} \wedge \mathrm{E3}).$$

By Lemmas 14, 15, 16, and 18, this is at least $1 - 3\times 10^{-5} - e^{-1/2^{12}} \ge .0002$.

4 Improvements

In this paper, we showed that n-user Binary Exponential Backoff is stable as long as the arrival rate is $O(n^{-(.75+\delta)})$ for any constant $\delta > 0$. A natural question is whether the protocol remains stable for higher arrival rates. In particular, it would be very interesting to know whether it is stable for some constant arrival rate.

Recall Foster's Theorem (Theorem 4) and Fayolle, Malyshev and Menshikov's generalisation of it (Theorem 5) from Section 3. Both theorems show that the relevant Markov chain is positive recurrent *if and only if* there is a potential function f satisfying the given conditions. Thus, if it turns out that binary exponential backoff is stable for higher arrival rates, the same proof technique could be used to prove the theorem.

On the other hand, finding an appropriate potential function might get increasingly difficult as the arrival rate gets higher. Furthermore, the number of cases that need to be considered may grow. Using our particular potential function, and our choice of cases to be considered,[5] we cannot prove stability for rates higher than about $n^{-.75}$. We suspect that our analysis would have to be improved substantially to show that the protocol is stable for any constant arrival rate. In particular, the analysis technique that we use in Case 2 seems too weak. After the preamble τ_0, the backoff counters are suitably high but still we do not show that a *constant factor* of the remaining steps have successes. Showing this (if indeed it is true!) would require a careful analysis of the situation immediately following each success, perhaps along the lines of the "capture" analysis in [7].

References

[1] D. Aldous, Ultimate instability of exponential back-off protocol for acknowledgement-based transmission control of random access communication channels, *IEEE Trans. Inf. Theory* **IT-33(2)** (1987) 219–233.

[2] G. Fayolle, V.A. Malyshev and M.V. Menshikov, *Topics in the Constructive Theory of Countable Markov Chains,* (Cambridge Univ. Press, 1995)

[3] L.A. Goldberg and P.D. MacKenzie, Analysis of practical backoff protocols for contention resolution with multiple servers, *Journal of Computer and Systems Sciences*, **58** (1999) 232–258.

[4] L.A. Goldberg, P.D. MacKenzie, M. Paterson and A. Srinivasan, Contention resolution with constant expected delay, Pre-print (1999) available at `http://www.dcs.warwick.ac.uk/~leslie/pub.html`. (Extends a paper by the first two authors in *Proc. of the Symposium on Foundations of Computer Science (IEEE)* 1997 and a paper by the second two authors in *Proc. of the Symposium on Foundations of Computer Science (IEEE)* 1995.)

[5] J. Goodman, A.G. Greenberg, N. Madras and P. March, Stability of binary exponential backoff, *J. of the ACM*, **35(3)** (1988) 579–602.

[5] For example, in Case 1, $g(0,1)$ is only negative if $\mu \geq \eta$ and $g(m,n)$ is only negative if $\mu + \eta \leq .5$. This forces $\eta \leq .25$.

[6] G.R. Grimmet and D.R. Stirzaker, *Probability and Random Processes, Second Edition.* (Oxford University Press, 1992)

[7] J. Håstad, T. Leighton and B. Rogoff, Analysis of backoff protocols for multiple access channels, *SIAM Journal on Computing* **25(4)** (1996) 740-774.

[8] F.P. Kelly, Stochastic models of computer communication systems, *J.R. Statist. Soc. B* **47(3)** (1985) 379–395.

[9] F.P. Kelly and I.M. MacPhee, The number of packets transmitted by collision detect random access schemes, *The Annals of Probability,* **15(4)** (1987) 1557–1568.

[10] R.M. Metcalfe and D.R. Boggs, Ethernet: Distributed packet switching for local computer networks. *Commun. ACM,* **19** (1976) 395–404.

[11] P. Raghavan and E. Upfal, Stochastic contention resolution with short delays, *SIAM Journal on Computing* **28** (1999) 709–719.

Appendix: Supplementary Calculations for Case 2

Here we show the inequality

$$\alpha n^{2-\eta-\mu}\lambda k - \alpha n^{2-\eta-\mu} W p + kn \leq -\epsilon k$$

holds when $\alpha p \geq 2^{13}$ and n is sufficiently large.

Case A: ($m \geq n^{1-\eta-\mu}$) In this case, $r = m$, $W = n^{\eta+\mu}\lceil \lg r\rceil 2^{-8}$, $v = n$, and $k = 4(r + v)\lceil \lg r\rceil$. Then, since $k \leq 8n\lceil \log m\rceil$, for large n,

$$\begin{aligned}
&\alpha n^{2-\eta-\mu}\lambda k - \alpha n^{2-\eta-\mu} W p + kn \\
&\leq \alpha n^{2-\eta-\mu}(\alpha' n^{1-\eta})^{-1}k - 2^{13} n^{2-\eta-\mu} W + kn \\
&\leq (\alpha/\alpha')n^{1-\mu}k - 2^{13}n^{2-\eta-\mu}n^{\eta+\mu}\lceil \log m\rceil 2^{-8} + kn \\
&\leq (\alpha/\alpha')n^{1-\mu}(4(m+n)\lceil \log m\rceil) - 2^5 n^2\lceil \log m\rceil + 4(m+n)\lceil \log m\rceil n \\
&\leq 8n^{2-\mu}\lceil \log m\rceil - 32n^2\lceil \log m\rceil + 8n^2\lceil \log m\rceil \\
&\leq -16n^2\lceil \log m\rceil \\
&\leq -2nk \\
&\leq -\epsilon k.
\end{aligned}$$

Case B: ($m < n^{1-\eta-\mu}$, $m' > n^{.4}$) In this case, $r = m'$, $W = \lceil \lg r\rceil 2^{-8}$, and $v = 2\lceil n^{1-\eta-\mu}\rceil$, and $k = 4(r+v)\lceil \lg r\rceil$. Note that by definition, $m' < n^{1-\eta-\mu}$. Then, since $k \leq 12\lceil n^{1-\eta-\mu}\rceil\lceil \log m'\rceil$, for large n,

$$\begin{aligned}
&\alpha n^{2-\eta-\mu}\lambda k - \alpha n^{2-\eta-\mu} W p + kn \\
&\leq \alpha n^{2-\eta-\mu}(\alpha' n^{1-\eta})^{-1}k - 2^{13} n^{2-\eta-\mu} W + kn \\
&\leq (\alpha/\alpha')n^{1-\mu}k - 2^{13}n^{2-\eta-\mu}\lceil \log m'\rceil 2^{-8} + kn \\
&\leq (\alpha/\alpha')n^{1-\mu}(4(m' + 2\lceil n^{1-\eta-\mu}\rceil)\lceil \log m'\rceil) - 2^5 n^{2-\eta-\mu}\lceil \log m'\rceil \\
&\quad +4(m + 2\lceil n^{1-\eta-\mu}\rceil)\lceil \log m'\rceil n \\
&\leq 12n^{1-\mu}\lceil n^{1-\eta-\mu}\rceil\lceil \log m'\rceil) - 2^5 n^{2-\eta-\mu}\lceil \log m'\rceil + 12n\lceil n^{1-\eta-\mu}\rceil\lceil \log m'\rceil \\
&\leq 12n^{1-\mu}\lceil n^{1-\eta-\mu}\rceil\lceil \log m'\rceil) - 32n^{2-\eta-\mu}\lceil \log m'\rceil + 13n^{2-\eta-\mu}\lceil \log m'\rceil \\
&\leq -18n^{2-\eta-\mu}\lceil \log m'\rceil \\
&\leq -nk \\
&\leq -\epsilon k.
\end{aligned}$$