Binary Exponential Backoff is Stable for High Arrival Rates

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

July 28, 1999

Abstract

Goodman, Greenberg, Madras and March gave a lower bound of $n^{-\Omega(\log n)}$ for the maximum arriaval 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 $O(n^{-.9})$.

Classification of topic: Algorithms and data structures (distributed algorithms)

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 [9]. 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 [6, 9] for details.

Håstad, Leighton and Rogoff [6] 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,

* 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 the EPSRC grant GR/L6098 and by ESPRIT Project RAND-II (Project 21726).

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

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 to the end of each queue independently with probability λ/n, where λ is the *arrival rate* of the system. 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 $X(t+1)$. One measure of the performance of the system is the expectation of the random variable T_{ret}, which is the number of steps required for the system to return to the start state $X(0) = ((0, \ldots, 0), (0, \ldots, 0))$. Håstad et al. [6] proved that if the arrival rate is too high, then the system is unstable, in the sense that the expected recurrence time is infinite.

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 system 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, as long as n is sufficiently large and $\lambda < \frac{1}{\alpha n^{.9}}$ then $E[T_{ret}]$ is finite for the n-user system.*

The point of Theorem 3 is to show that n-user Binary Exponential Backoff is stable for arrival rates which grow faster asymptotically than $1/n$. That is, the purpose of the result is to show that, for positive constants α and η, $\lambda < \frac{1}{\alpha n^{1-\eta}}$ guarantees stability. We have chosen $\eta = .1$ for concreteness. We believe that the same methods could be used for slightly larger values of η, but an interesting (and difficult) question raised by this work is whether the same result would be true for $\eta = 1$. That is, is there a constant α such that the n-user system is stable whenever $\lambda < \frac{1}{\alpha}$?

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 [7, 8] showed this for $\lambda < \ln 2$ and Aldous [1] showed that it holds for all positive λ.

Finally, we mention that, 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 [10] 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 [6] to show that, when λ is sufficiently low, the expected average message waiting time is finite.

3 The stability proof

In order to prove Theorem 3, let $\lambda = \frac{1}{\alpha' n^{.9}}$, where $\alpha' \geq \alpha$. We will now define the relevant potential function. Let $f(X(t))$ be the following function of the state just before step t.

$$f(X(t)) = \alpha n^{1.8} \sum_{i=1}^{n} q_i(t) + \sum_{i=1}^{n} 2^{b_i(t)}.$$

We will use the following generalisation of Foster's theorem [2]. Note that the Markov chain X satisfies the initial conditions of the theorem. That is, it is time-homogeneous, irreducible, and aperiodic and has a countable state space.

Theorem 4 (Foster; 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.*

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

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

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 $b_i(t) < .8 \lg n + 1$. 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 2 is satisfied. For every state $\rho \notin C$, we will define $k(\rho)$ in such a way that Equation 1 is also satisfied. We give the details in three cases.

3.1 Case 1: $m'(X(t)) = 0$ and $m(X(t)) < n^{.8}$.

For every state ρ such that $m'(\rho) = 0$ and $m(\rho) < n^{.8}$ we define $k(\rho) = 1$. We wish to show that, if $\rho \neq ((0,\ldots,0),(0,\ldots,0))$ and $X(t) = \rho$, then $E[f(X(t+1) - f(X(t))] \leq -\epsilon$. Our general approach is the same as the approach used in the proof of Lemma 5.7 of [6]. 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

$$E[f(X(t+1) - f(X(t))] = \alpha n^{1.8} \sum_{i=1}^n (E[I_{a,i,t}] - E[I_{s,i,t}]) + \sum_{i=1}^n \left(E[2^{b_i(t+1)}] - 2^{b_i(t)}\right),$$

$$= \alpha n^{1.8}\lambda - \alpha n^{1.8} ST + \sum_{i=1}^n \left(2^{b_i(t)}\sigma_i - (2^{b_i(t)} - 1)\pi_i\right), \quad (3)$$

$$= \alpha n^{1.8}\lambda - \alpha n^{1.8} ST + \sum_{i=1}^n \left(2^{b_i(t)} p_i (1 - \frac{T}{1-p_i}) - (2^{b_i(t)} - 1) p_i \frac{T}{1-p_i}\right),$$

$$= \alpha n^{1.8}\lambda - \alpha n^{1.8} ST + \sum_{i=1}^{\ell} (1 - \frac{T}{1-p_i}) + \sum_{i=\ell+1}^{n} \frac{\lambda}{n}(1 - \frac{T}{1-p_i}) - \ell T,$$

$$= \alpha n^{1.8}\lambda - \alpha n^{1.8} ST + \ell - \ell T + \frac{(n-\ell)\lambda}{n} - T\left(\sum_{i=1}^{\ell} \frac{1}{1-p_i} + \sum_{i=\ell+1}^{n} \frac{p_i}{1-p_i}\right),$$

$$= \alpha n^{1.8}\lambda - \alpha n^{1.8} ST + \ell - \ell T + \frac{(n-\ell)\lambda}{n} - ST - \ell T,$$

$$= \alpha n^{1.8}\lambda + \ell + \frac{(n-\ell)\lambda}{n} - T((\alpha n^{1.8} + 1)S + 2\ell), \quad (4)$$

where σ_i in Equality 3 denotes the probability that user i collides at step t and π_i denotes the probability that user i sends successfully at step t. We now find lower bounds for S and T. First,

$$\begin{aligned}
S &= \sum_{i=1}^{n} \frac{p_i}{1-p_i} \\
&= \sum_{i=1}^{\ell} \left(\frac{2^{-b_i(t)}}{1-2^{-b_i(t)}} \right) + \frac{\lambda(n-\ell)}{n-\lambda} \\
&\geq \sum_{i=1}^{m} \left(\frac{1}{\beta n - 1} \right) + \frac{\lambda(n-\ell)}{n-\lambda} \\
&= \frac{m}{\beta n - 1} + \frac{\lambda(n-\ell)}{n-\lambda}. \qquad (5)
\end{aligned}$$

Next,

$$\begin{aligned}
T &= \prod_{i=1}^{n} (1-p_i) \\
&\geq (1-\frac{1}{2n^{.8}})^{m} (1-\frac{1}{\beta n})^{\ell-m} (1-\frac{\lambda}{n})^{n-\ell} \\
&\geq 1 - \frac{m}{2n^{.8}} - \frac{\ell-m}{\beta n} - \frac{\lambda(n-\ell)}{n} \qquad (6)
\end{aligned}$$

Combining Equations 4, 5 and 6, we get the following equation.

$$E[f(X(t+1)-f(X(t))] \leq \alpha n^{1.8}\lambda + \ell + \frac{(n-\ell)\lambda}{n} - \left(1 - \frac{m}{2n^{.8}} - \frac{\ell-m}{\beta n} - \frac{\lambda(n-\ell)}{n}\right)\left((\alpha n^{1.8}+1)\left(\frac{m}{\beta n-1} + \frac{\lambda(n-\ell)}{n-\lambda}\right) + 2\ell\right). \qquad (7)$$

We will let $g(m,\ell)$ be the quantity in Equation 7 plus ϵ and we will show that $g(m,\ell)$ is negative for all values of $0 \leq m < n^{.8}$ 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. We include them in Appendix A for completeness.

3.2 Case 2: $m(X(t)) \geq n^{.8}$ or $m'(X(t)) > n^{.4}$.

For every state ρ such that $m(\rho) \geq n^{.8}$ 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 $m \geq n^{.8}$ then we will define $r = m$, $W = m^{1/4}\lceil \lg r \rceil 2^{-8}$, $A = W$, $b = \lg \beta + \lg n$ and $v = n$. Otherwise, we will define $r = m'$, $W = \lceil \lg r \rceil 2^{-8}$, $A = 0$, $b = .8 \lg n + 1$, and $v = 2\lceil n^{.8} \rceil$. In either case, we will define $k = 4(r+v)\lceil \lg r \rceil$. Let τ be the set of 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

$$\begin{aligned} E[f(X(t+k) - f(X(t))] &\leq \alpha n^{1.8}\lambda k - \alpha n^{1.8} E[\mathcal{S}] + \sum_{i=1}^{n} \sum_{t'=t+1}^{t+k} E[2^{b_i(t')} - 2^{b_i(t'-1)}] \\ &\leq \alpha n^{1.8}\lambda k - \alpha n^{1.8} W p + kn \\ &\leq -\epsilon k, \end{aligned}$$

where the final inequality holds as long as $\alpha p \geq 2^{13}$ and n is sufficiently big. 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 5 *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 steps $[t, \ldots, t + \delta j \lceil \lg j \rceil - 1]$, 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 $B = \lceil W \rceil + \lceil A \rceil$, $k' = 4r\lceil \lg r \rceil$, $k'' = 4B\lceil \lg B \rceil$ and $\tau_0 = \{t, \ldots, t + k' - 1\}$. Let $\tau'(i)$ be the set of all $t' \in \tau$ such that $b_i(t') = 0$ and either (1) $q_i(t') > 0$ or (2) there is an arrival at user i at t'. Let τ_2 be the set of all $t' \in \tau$ such that $|\{(t'', i) \mid t'' \in \tau'(i) \text{ and } t'' < t'\}| \geq B$. Finally, let τ_1 be the set of all $t' \in \tau - \tau_0 - \tau_2$ such that, for some i, $\tau'(i) \cap [t' - k'' + 1, t'] \neq \emptyset$. 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. Every station with $q_i(t) > 0$ and $b_i(t) < b$ has $b_i(t') \leq b + \lg(r)/2 + 3$ for all $t' \in \tau$.

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

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

Lemma 6 *If n is sufficiently large then* $\Pr(\overline{E1}) \leq 10^{-5}$.

Proof: The expected number of arrivals in τ is λk. If $m \geq n^{.8}$, then $A = m^{1/4}\lceil \lg r \rceil 2^{-8} \geq 2\lambda k$. 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, $\Pr(\text{E1}) \geq (1 - \lambda/n)^{nk} \geq 1 - \lambda k \geq 1 - 10^{-5}$. □

Lemma 7 *If n is sufficiently large then* $\Pr(\overline{E2}) \leq 10^{-5}$.

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

Lemma 8 *If n is sufficiently large then* $\Pr(\overline{E3}) \leq 10^{-5}$.

Proof: Let $y = \left\lceil \frac{\lceil \lg r \rceil}{4} \right\rceil$. Note that $2y \leq \frac{\lg r}{2} + 3$. Suppose that user i has $b_i(t') > b + 2y$. Then this user sent when its backoff counter was $\lceil b + z \rceil$ for all $z \in \{y, \ldots, 2y - 1\}$. The probability of such a send on any particular step is at most $\frac{1}{2^b 2^y}$. Thus, the probability that it makes all y of the sends is at most

$$\binom{k}{y}\left(\frac{1}{2^b 2^y}\right)^y \leq \left(\frac{ke}{2^b y 2^y}\right)^y \leq 10^{-5}/r.$$

Thus, the probability that any of the r users obtains such a big backoff counter is at most 10^{-5}. □

Lemma 9 *If n is sufficiently large then* $\Pr(\overline{E4}) \leq 10^{-5}$.

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

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

Lemma 10 *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 $b_i(t+z) \leq b + \lg(r)/2 + 3$, the probability that user i succeeds at step $t + z$ is at least $\frac{1}{2^{10} 2^b r^{1/2}}$.*

Proof: The conditions in the lemma imply the following.

- There are no users j with $b_j(t+z) < \lceil \lg B \rceil$.
- There are at most B users j with $b_j(t+z) < \lceil \lg r \rceil$.

- There are at most $r+B$ users j with $b_j(t+z) < b$.
- There are at most $m+B$ users j with $b_j(t+z) < \lg\beta + \lg n$.

Thus, the probability that user i succeeds is at least

$$\begin{aligned}
& 2^{-(b+\lg(r)/2+3)}\left(1-\frac{1}{B}\right)^B\left(1-\frac{1}{r}\right)^r\left(1-\frac{1}{2^b}\right)^{m-r}\left(1-\frac{1}{\beta n}\right)^{n-m-B} \\
\geq & \frac{1}{2^b r^{1/2} 2^3}\frac{1}{4}\frac{1}{4}\frac{1}{4}\left(1-\frac{n-m-B}{\beta n}\right) \\
\geq & \frac{1}{2^{10} 2^b r^{1/2}}.
\end{aligned}$$

□

Corollary 11 *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 $\frac{r-B}{2^{10}2^b r^{1/2}} \geq \frac{1}{2^{13}n^{.6}}$.*

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, each of these has $b_i(t+z) \leq b + \lg(r)/2 + 3$. 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 $t+z$. □

Lemma 12 *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 11 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 $\frac{1}{2^{13}n^{.6}}$. 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 $\exp(-\frac{v\lceil\lg r\rceil}{2^{17}n^{.6}}) \leq 10^{-5}$. □

We conclude Case 2 by observing that p is at least $1-\Pr(\overline{E1})-\Pr(\overline{E2})-\Pr(\overline{E3})-\Pr(\overline{E4})-\Pr(\mathcal{S} < W \mid E1 \wedge E2 \wedge E3 \wedge E4)$. By Lemmas 6, 7, 8, 9, and 12, this is at least $1-5\times 10^{-5}$.

3.3 Case 3: $0 < m'(X(t)) \leq n^{.4}$ and $m(X(t)) < n^{.8}$.

For every state ρ such that $0 < m'(\rho) \leq n^{.4}$ and $m(\rho) < n^{.8}$, we will define $k = 32m'(\rho)\lceil\lg m'(\rho)\rceil + \lceil n^{.8}\rceil$. We will show that, if $X(t)=\rho$, then $E[f(X(t+k)-f(X(t))] \leq -\epsilon k$. 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, $E[f(X(t+k)-f(X(t))] \leq \alpha n^{1.8}\lambda k - \alpha n^{1.8} p + kn$, 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 $b_\gamma(t) < .8\lg n + 1$. 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 $b_i(t) < .8\lg n + 1$ either sends successfully during τ_0 or has $b_i(t + k') \geq \lceil \lg m' \rceil$.

E3. $b_\gamma(t') < .8\lg n + 7$ for all $t' \in \tau$.

Lemma 13 *If n is sufficiently large then $Pr(\overline{E1}) \leq 10^{-5}$.*

Proof: As in the proof of Lemma 6,

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

□

Lemma 14 $Pr(\overline{E2}) \leq 10^{-5}$.

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

$$Pr(\overline{E3}) \leq m' \cdot \frac{\lceil \lg m' \rceil}{(m')^{16/\ln(2)}} \leq 10^{-5}. \tag{8}$$

□

Lemma 15 *If n is sufficiently large then $Pr(\overline{E3}) \leq 10^{-5}$.*

Proof: Let $y = 6$, and suppose that user γ sends with backoff $b_\gamma = \lceil .8\lg n + r \rceil$ for $r \in \{1, \ldots, 6\}$. The probability of this happening is

$$\begin{aligned} Pr(\overline{E3}) &\leq \binom{k}{6} \prod_{r=1}^{6} 2^{-\lceil .8\lg n \rceil - r} \\ &\leq \left(\frac{ke}{6}\right)^6 \left(\frac{1}{n^{.8}}\right) 2^{-\sum_{r=1}^{6} r} \\ &\leq \left(\frac{2en^{.8}}{6n^{.8}2^3}\right)^6 \\ &\leq 10^{-5}. \end{aligned}$$

□

Lemma 16 *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 $\frac{1}{2^{12}n^{.8}}$.*

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

- $q_\gamma(t+z) > 0$ and $b_\gamma(t+z) < .8\lg n + 7$.
- There are no users j with $b_j(t+z) < \lceil \lg m' \rceil$.
- There are at most m' users j with $b_j(t+z) < .8\lg n + 1$.

- 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

$$\begin{aligned} & 2^{-(.8\lg n+7)}\left(1-\frac{1}{m'}\right)^{m'}\left(1-\frac{1}{2n^{.8}}\right)^{m-m'}\left(1-\frac{1}{\beta n}\right)^{n-m} \\ \geq\ & \frac{1}{2^7 n^{.8}}\frac{1}{4}\frac{1}{4}\frac{1}{2} \\ \geq\ & \frac{1}{2^{12}n^{.8}}. \end{aligned}$$

□

Lemma 17 *If n is sufficiently large then $Pr(\mathcal{S} < 1 \mid E1 \wedge E2 \wedge E3) \leq e^{-1/2^{12}}$.*

Proof: Lemma 16 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 $\frac{1}{2^{12}n^{.8}}$. Since $|\tau - \tau_0| \geq n^{.8}$, this probability is at most

$$\left(1-\frac{1}{2^{12}n^{.8}}\right)^{n^{.8}} \leq e^{-1/2^{12}}.$$

□

We conclude Case 3 by observing that p is at least $1-\Pr(\overline{\mathrm{E1}})-\Pr(\overline{\mathrm{E2}})-\Pr(\overline{\mathrm{E3}})-\Pr(\mathcal{S} < 1 \mid \mathrm{E1}\wedge\mathrm{E2}\wedge\mathrm{E3})$. By Lemmas 13, 14, 15, and 17, this is at least $1-3\times 10^{-5}-e^{-1/2^{12}} \geq .0002$.

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.

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

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

[8] 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.

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

[10] P. Raghavan and E. Upfal, Contention resolution with bounded delay, *Proc. of the ACM Symposium on the Theory of Computing* **24** (1995) 229–237.

4 Appendix A: The Calculations for Case 1.

1. $g(m,m)$ **is negative:** $g(m,m) \times 2\alpha' n^{1.9}(\beta n - 1)(\alpha' n^{1.9} - 1)$ is equal to the following.

$$
\begin{aligned}
- \quad & 2m - 4m^2 + 2n + 6mn + 2\beta mn + 6\beta m^2 n + 2\alpha' m^2 n^{1.1} + 2\alpha' \epsilon n^{1.9} + 2\alpha' m^2 n^{1.9} \\
- \quad & 2n^2 - 2\beta n^2 - 8\beta mn^2 - \alpha' mn^{2.1} - 3\alpha' \beta m^2 n^{2.1} + 2\alpha n^{2.8} + 2\alpha mn^{2.8} + 2\alpha\beta m^2 n^{2.8} \\
- \quad & 2\alpha' \beta\epsilon\, n^{2.9} - 2\alpha' mn^{2.9} + 2\alpha' \beta mn^{2.9} - 4\alpha' \beta m^2 n^{2.9} + 2\beta n^3 - \alpha'^2 m^2 n^3 + \alpha' \beta mn^{3.1} \\
- \quad & 2\alpha\alpha' m^2 n^{3.7} - 2\alpha n^{3.8} - 2\alpha\beta n^{3.8} - 2\alpha'^2 \epsilon n^{3.8} - 4\ \alpha\beta mn^{3.8} - \alpha\alpha' mn^{3.9} + 4\alpha' \beta mn^{3.9} \\
- \quad & \alpha\alpha' \beta m^2 n^{3.9} + 2\alpha'^2 \beta m^2 n^4 + 2\alpha\alpha' mn^{4.7} + 2\ \alpha\alpha' \beta mn^{4.7} + 2\alpha\beta n^{4.8} + 2\alpha'^2 \beta\epsilon n^{4.8} \\
- \quad & 2\alpha'^2 \beta mn^{4.8} + \alpha\alpha'^2 m^2 n^{4.8} + \alpha\alpha' \beta mn^{4.9} - 2\alpha\alpha'^2 mn^{5.6}
\end{aligned}
$$

The dominant term is $-2\alpha\alpha'^2 mn^{5.6}$. Note that there is a positive term ($\alpha\alpha'^2 m^2 n^{4.8}$) which could be half this big if m is as big as $n^{.8}$ (the upper bound for Case 1), but all other terms are asymptotically smaller.

2. $g(m,n)$ **is negative:** $g(m,n) \times 2\alpha' \beta n(\beta n - 1)$ is equal to the following.

$$
\begin{aligned}
- \quad & 2\alpha' m^2 + \alpha' \beta m^2 n^{.2} - 2\alpha' \beta\epsilon n + 6\alpha' mn - 2\alpha' \beta mn - 2\alpha' \beta mn^{1.2} \\
- \quad & 2\alpha\alpha' m^2 n^{1.8} - 2\alpha\beta n^{1.9} - 4\alpha' n^2 + 2\alpha' \beta n^2 + 2\alpha' \beta^2 \epsilon n^2 \\
- \quad & 4\alpha' \beta mn^2 + \alpha\alpha' \beta m^2 n^2 + 2\alpha' \beta^2 mn^{2.2} + 2\alpha\alpha' mn^{2.8} - 2\alpha\alpha' \beta mn^{2.8} + 2\alpha\beta^2 n^{2.9} \\
+ \quad & 4\alpha' \beta n^3 - 2\alpha' \beta^2 n^3
\end{aligned}
$$

Since $\beta > 2$, the term $-2\alpha' \beta^2 n^3$ dominates $+4\alpha' \beta n^3$. For the same reason, the term $-2\alpha\alpha' \beta mn^{2.8}$ dominates the two terms $+2\alpha\alpha' mn^{2.8}$ and $+\alpha\alpha' \beta m^2 n^2$. The other terms are asymptotically smaller.

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

$$\frac{\partial^2}{\partial \ell^2} g(m,\ell) = 2\left(\frac{1}{\beta n} - \frac{\lambda}{n}\right)\left(2 - \frac{(\alpha n^{1.8} + 1)\lambda}{n - \lambda}\right).$$

1'. $g(0,1)$ **is negative:** $g(0,1) \times \alpha' \beta n^{1.9}(\alpha' n^{1.9} - 1)$ is equal to the following.

$$
\begin{aligned}
+ \quad & 4\beta - 3\alpha' n^{.9} - 5\beta n + \alpha\beta n^{1.8} + \alpha' n^{1.9} - \alpha' \beta m^{1/9} - \alpha' \beta\epsilon n^{1.9} \\
+ \quad & \beta n^2 - \alpha\alpha' n^{2.7} + 2\alpha'^2 n^{2.8} - 3\alpha\beta n^{2.8} + 2\alpha' \beta n^{2.9} + \alpha\alpha' n^{3.7} \\
+ \quad & \alpha\alpha' \beta n^{3.7} + \alpha\beta n^{3.8} - \alpha'^2 \beta n^{3.8} + \alpha'^2 \beta\epsilon n^{3.8}
\end{aligned}
$$

Since $\alpha'(1-\epsilon) \geq \alpha(1-\epsilon) > 1$, the term $-\alpha'^2 \beta(1-\epsilon)n^{3.8}$ dominates the term $+\alpha\beta n^{3.8}$. The other terms are asymptotically smaller.