SIAM J. COMPUT.
Vol. 28, No. 2, pp. 709–719

STOCHASTIC CONTENTION RESOLUTION WITH SHORT DELAYS*

PRABHAKAR RAGHAVAN[†] AND ELI UPFAL[‡]

Abstract. We study contention resolution protocols under a stochastic model of continuous request generation from a set of contenders. The performance of such a protocol is characterized by two parameters: the maximum arrival rate for which the protocol is stable and the expected delay of a request from arrival to service.

Known solutions are either unstable for any constant injection rate or have at least polynomial (in the number of contenders) expected delay. Our main contribution is a protocol that is stable for a constant injection rate, while achieving logarithmic expected delay. We extend our results to the case of multiple servers, with each request being targeted for a specific server. This is related to the *optically connected parallel computer* (or *OCPC*) model. Finally, we prove a lower bound showing that long delays are inevitable in a class of protocols including backoff-style protocols, if the arrival rate is large enough (but still smaller than 1).

Key words. contention resolution, randomized algorithms, stochastic analysis

AMS subject classifications. 68Q25, 68Q75, 68M20, 60K30

PII. S0097539795285333

1. Introduction.

The subject of this paper is the stochastic analysis of protocols for contention resolution. The most concrete setting of this problem is that of multiple access channels, and so the remainder of the paper will use the terminology of this application. Naturally, our analyses are not specific to this setting and apply whenever we have several contenders requesting service from shared resources.

There are n *senders* and m *receivers*. At each of a series of time steps, one or more senders may generate a *packet*. A packet when generated has a *destination*: a unique receiver to which it must be delivered. Any sender may attempt to send a packet to any receiver at any step, but a receiver may only receive one packet in a step. If a receiver is sent more than one packet in a step (a *collision*), all packets sent to that receiver are lost and the senders are notified of the loss. The senders must then try to send these packets again at a future step. There is no explicit communication between the senders for coordinating the transmissions the only information that senders have is the packet(s) they have waiting for transmission and the history of losses. A packet can only be transmitted directly from its sender to its receiver; intermediate hops are disallowed.

The case $m = 1$ is a classical instance of sharing a common resource such as a bus or an Ethernet channel (the shared bus is modeled by the single "receiver"). The binary exponential backoff Ethernet protocol [9] is the solution used most commonly in practice here. The case $m = n$ has received much attention recently under the name *optically connected parallel computer*, or OCPC [8]. However, work on the OCPC model has been restricted to studies in which each sender begins with at most h

---

[*] Received by the editors April 24, 1995; accepted for publication (in revised form) March 28, 1997; published electronically August 4, 1998.

http://www.siam.org/journals/sicomp/28-2/28533.html

[†] IBM Almaden Research Center, 1650 Harry Rd., San Jose, CA 95120 (pragh@almaden.ibm.com). A portion of this work was done at the IBM T.J. Watson Research Center, Yorktown Heights, NY 10598.

[‡] The Weizmann Institute of Science, Israel (eli@wisdom.weizmann.ac.il), and IBM Almaden Research Center, 1650 Harry Rd., San Jose, CA 95120. Work at the Weizmann Institute of Science was supported in part by the Norman D. Cohen Professorial Chair of Computer Science, a MINERVA grant, and a grant from the Israeli Academy of Science.

packets for some positive integer h, with h or fewer packets bound for each receiver. In our work, we are interested in the more realistic setting in which packets are generated continuously. To this end, we adopt a stochastic model of packet generation.

1.1. The model. We assume that time is partitioned into intervals of equal length called *steps*. The injection distribution is characterized by an $n \times m$ matrix $\Lambda = (\lambda_{ij})$, in which $\sum_j \lambda_{ij} \leq \lambda$ for all i and $\sum_i \lambda_{ij} \leq \lambda$ for all j. At each step sender i generates a packet bound for receiver j with probability λ_{ij} independently of other senders and time steps. Since a sender can send no more than one packet per step, and a receiver can receive no more than one packet per step, we require $\lambda \leq 1$. For the special case $m = 1$, we use simplified notation: we assume that at each step sender i generates a packet with probability λ_i, independently of other steps and senders. Note that a sender generates at most one packet at a step, but several senders may generate a packet for a receiver. The reader is referred to the paper by Håstad, Leighton, and Rogoff [6] for a detailed account of the relation between these assumptions and reality.

Each sender uses a *protocol* to decide when to send a packet and what to do in the event of a collision. Informally, a protocol is an automaton that uses its state to remember the past. At each step, based on its state, the packets pending transmission, and any new packet it generates, it must decide which (if any) packet to transmit. For example, the binary exponential backoff protocol used in the Ethernet is the following: store arriving packets in a queue. When a packet reaches the head of the queue for the first time, transmit it. If there is a collision, retransmit the packet after T steps where T is selected randomly from $\{1, 2, 3, \ldots, 2^{\min\{10,b\}}\}$, where b is the number of times the packet has been involved in a collision.

Of primary interest is whether a protocol can sustain packet arrivals at some rate without *instability:* a protocol is unstable for a given arrival rate if the number of packets pending transmission grows unboundedly with time.[1] For stable protocols, two quantities are of interest: (1) the *maximum arrival rate* λ that can be sustained stably, and (2) the *delay*, defined to be the maximum over all senders of the expected number of steps from the generation of a packet to its delivery, in the steady state. Delay is of particular importance in high-speed communications applications such as video and ATM networks. (We actually prove a stronger result, a bound on the expected delay of any packet arriving after a polynomial number of steps from the start of the process.)

In our upper bounds, we do not assume that the sender knows the injection rates of other senders (i.e., the matrix Λ); we only assume that senders have some upper bound on the total number of active senders in the system. This bound does not need to be very accurate, as only the logarithm of this bound figures in the delay of our protocols. Since the total number of senders in the system is dictated by the hardware, it is reasonable to assume that this number cannot change very fast and that each sender has a good estimate of it. Our analysis holds even if the matrix Λ changes at every time step, subject to the constraints on row and column sums.

Throughout this paper, $\mathbf{Pr}[\mathcal{E}]$ denotes the probability of an event $\mathcal{E}$.

1.2. Related work. Most previous analysis focused on backoff protocols. The binary exponential backoff protocol used in the Ethernet was proposed by Metcalfe

---

[1] Formally let B_t denote the number of packets generated before time t that were not delivered until time t. Let F_t denote the distribution of B_t. A protocol is stable if the sequence $\{F_t \mid t \geq 1\}$ converges in distribution to a limit distribution F that is independent of t.

and Boggs [9]. Aldous [1] showed that for any positive constant λ, binary exponential backoff is unstable if the number of senders is infinite. Kelly [7] showed that any polynomial backoff protocol is unstable for infinitely many senders. Håstad, Leighton, and Rogoff [6] studied systems with a finite number of senders. They showed that binary exponential backoff is unstable for λ slightly larger than 0.567 even for a system with a finite number of senders, and that polynomial backoff protocol is stable for any $\lambda < 1$ and for any finite number of senders. Goldberg and MacKenzie [5] analyzed the backoff protocol for the multiple servers setting, showing that any superlinear polynomial backoff protocol is stable for any $\lambda < 1$. Our current work is motivated by the lower bound in Håstad, Leighton, and Rogoff [6] showing that long delays are inevitable in backoff-style protocols: they showed that the delay of any stable exponential or polynomial backoff protocol is at least polynomial in the total number of contenders. Following our work, Paterson and Srinivasan [11] recently gave a protocol with $O(1)$ expected delay assuming that the initial clock times of the senders are within a known bound of each other.

1.3. Our results.

We follow the lead of [1, 6] here and make no unproven assumptions about the independence of the state of the system from one time step to the next, or between senders (many analyses in the queuing-theory literature do make such assumptions). We assume only that the generation of packets is independent between time steps and senders. Even for this case, complex dependencies arise between the transmissions at different senders and time steps. We focus on protocols with short packet delay. We first consider the case $m = 1$. Our main result (Theorem 1) is a protocol that ensures delay logarithmic in n, provided the arrival rate is no more than a fixed constant λ'. To our knowledge, this is the first protocol with sublinear delay (under an exact analysis) that is stable for a constant injection rate. Turning to the case $m > 1$, we present a protocol that achieves logarithmic delay provided that arrival rate at each sender and for each receiver is at most a fixed constant λ' (Theorem 2). Finally, we show (Theorem 3) that if every sender uses a class of protocols including backoff protocols, there is a fixed constant $\lambda_0 < 1$ such that if $\lambda > \lambda_0$, the delay must be $\Omega(n)$. Thus, in this class of protocols one cannot achieve full throughput and small delay simultaneously (whereas full throughput alone can be achieved by the polynomial backoff protocol that belongs to that class [6]).

2. Multiple access to one channel.

In this section we consider the case $m = 1$. We show that there exists a positive constant λ_0 and a contention resolution protocol that is stable for any $0 < \lambda \leq \lambda_0$, with delay $O(\log n)$.

2.1. The protocol.

Each sender in our protocol (Fig. 1) has a *transmission buffer* of size $O(\log n)$ and a *queue*. Packets awaiting transmission are stored either in the buffer or in the queue. Throughout the execution of the protocol a sender is in one of two states: a *normal state* or a *reset state*. Note that in the protocol, a transmission attempt may fail due to collision. The constants α and μ used in the protocol are fixed in the proof of Theorem 1 below.

Relation to practical protocols. The use of the queue and especially the reset state may appear to be somewhat artificial. We use these devices to cater to catastrophic events (e.g., every sender generates a packet in every one of n^{10} consecutive steps) that occur with extremely low but positive probabilities. In practice, such catastrophic events are handled by *dropping packets*: some packets are permanently erased from the system during such rare events. Rather than drop packets we invoke the emergency mechanism involving the queues and the reset state, while proving that the chance of resorting to these measures is extremely small.

```
$Count\_attempts(s)$ keeps a count of the number of times $s$ tried to transmit a packet
from its buffer in the $4\mu n \log n$ most recent steps.

$Failure\_counts(s)$ stores the failure rates in transmission attempts of packets from
the buffer of $s$ in the most recent $\mu \log n$ attempts.

$Random\_number()$ is a function that returns a random number uniformly chosen
in the range $[0, 1]$, independent of the outcomes of previous calls to the function.

While in the normal state repeat:
    1. Place new packets in the buffer.
    2. Let $X$ denote the number of packets in the buffer.
       If $Random\_number() \leq X/8\alpha \log n$ then
           (a) Try to transmit a random packet from the buffer.
           (b) Update $Count\_attempts(s)$ and $Failure\_counts(s)$.
       Else
       If $Random\_number() \geq 1 - 1/n^2$ then transmit the packet at the head of
       the queue.
    3. If ($Count\_attempts(s) \geq \mu \log n$ and $Failure\_counts(s) > 5/8$), or If
       ($X > 2\alpha \log n$) then
           (a) Move all packets in the buffer to the end of the queue.
           (b) Switch to the reset state for $4\mu n^2 \log n + \gamma \log n$ steps.
While in the reset state repeat:
    1. Append any new packets to the queue.
    2. If $Random\_number() \leq 1/n^2$ then transmit the packet at the head of the
       queue.
```

FIG. 1. *Communication protocol for sender s.*

2.2. Analysis of the protocol.

The performance of the protocol is summarized in the following theorem.

THEOREM 1. *There is a fixed constant $\lambda_0 > 0$, such that for any $\lambda \leq \lambda_0$ the above protocol is stable and the expected delay of each packet is $O(\log n)$.*

Proof. To simplify the analysis we assume, without loss of generality, that at each step each sender tries to transmit a packet from the queue with probability $1/n^2$ even if its queue is empty (in which case it tries to transmit an empty message). This assumption makes the process of transmitting from the buffers completely independent of the sizes of the queues at the beginning of that step.

Consider the n-vector of nonnegative integer whose ith component is the number of packets in the buffer of the ith sender. This vector defines a finite positive recurrent aperiodic Markov chain, which thus has a stationary distribution. Let X_t be a random variable counting the total number of packets at time t in all the buffers. By the condition tested in step 3 of the protocol, $X_t \leq 2\alpha n \log n + n$ for all t. The crux of the proof is to show that most of the time $X_t \leq 2\alpha \log n$, which guarantees short delays.

Throughout the analysis we use the following versions of the Chernoff bound [10]: let Z be the number of successes in k independent Bernoulli trials with probability p for success in each trial, then for $0 \leq \delta \leq 1$ $\mathbf{Pr}\{Z \leq (1-\delta)pk\} \leq e^{-\delta^2 pk/2}$, and $\mathbf{Pr}\{Z \geq (1+\delta)pk\} \leq e^{-\delta^2 pk/3}$. For $\delta > 1$ $\mathbf{Pr}\{Z \geq (1+\delta)pk\} \leq e^{-\delta \ln(1+\delta)pk}$. □

LEMMA 1. *Let* $T = 4\mu n^2 \log n + \gamma \log n$ (*the constant* γ *is determined in the proof of Claim* 3). *For any* $x \leq 2\alpha n \log n + n$ *assumed by* X_{t-T},

$$\mathbf{Pr}[X_t > 2\alpha \log n \mid X_{t-T} = x] \leq 1/n^{10}.$$

Proof. To prove the lemma, we begin with the following claim.

CLAIM 1. *If* $X_{t-T} > 6\alpha \log n$, *then with probability* $1 - n^{-11}$ *there is a step* $\tau_1 \in [t - T, t - \gamma \log n]$ *such that* $X_{\tau_1} \leq 6\alpha \log n$.

Proof. We show that as long as the total number of packets in buffers exceeds $6\alpha \log n$, then in each interval of $4\mu n \log n$ steps at least one sender is very likely to switch to the reset state with high probability (and stays in that state for $4\mu n^2 \log n$ steps).

Consider the interval $[\tau_0, \tau_0 + 4\mu n \log n - 1]$ consisting of $4\mu n \log n$ steps. Let z_τ, $\tau \in [\tau_0, \tau_0 + 4\mu n \log n - 1]$ be a random variable defined as follows. If $X_\tau > 6\alpha \log n$, then z_τ equals the number of transmission attempts at this step; else $z_\tau = n$. Clearly $[z_\tau \mid z_{\tau_0}, z_{\tau_0+1}, \ldots, z_{\tau-1}]$ is stochastically lower-bounded[2] by a binomial distribution with expectation $(6\alpha \log n)/(8\alpha \log n)$, and $\sum_{\tau=\tau_0}^{\tau_0+4\mu n \log n - 1} z_\tau$ is stochastically (lower-) bounded by a binomial distribution with expectation

$$\frac{(4\mu n \log n)(6\alpha \log n)}{8\alpha \log n} = 3\mu n \log n. \tag{1}$$

Thus, by the Chernoff bound, with probability at least

$$1 - e^{-\frac{2}{3}\mu n \log n}, \tag{2}$$

either $X_\tau \leq 6\alpha \log n$ for some τ in this interval or there were at least $\mu n \log n$ attempts to transmit packets from buffers and at least one sender was involved in $\mu \log n$ or more attempts.

If there were at least $6\alpha \log n$ packets in buffers in a given step, then the success probability of an attempt at that step is at most

$$\left(1 - \frac{2\alpha \log n}{8\alpha \log n}\right)^2 \leq 9/16. \tag{3}$$

(Since no sender has more than $2\alpha \log n$ packets in buffers, the "best" probability is when $6\alpha \log n$ packets are distributed equally among three senders.)

For a given sender s and given steps τ and τ' such that $\tau \geq \tau' \geq \tau_0$, let $y_\tau^{\tau',s}$ be a random variable defined as follows: $y_\tau^{\tau',s} = 1$ if sender s successfully transmitted a packet at time τ, $X_\tau > 6\alpha \log n$, and s had less than $\mu \log n$ transition attempts in the interval $[\tau', \tau]$; else $y_\tau^{\tau',s} = 0$. Let $\mathcal{H}_\tau$ describe the state of the system at all times before step τ. Clearly

$$\mathbf{Pr}[Y_\tau^{\tau',s} \mid \mathcal{H}_\tau] \leq \frac{9}{16},$$

and $\mathbf{Pr}[\sum_{\tau=\tau'}^{\tau_0+4\mu n \log n - 1} y_\tau^{\tau',s} \geq \frac{5}{8}\mu \log n]$ is upper-bounded by the probability that a binomial random variable with parameters $B(\mu \log n, \frac{9}{16})$ is at least $\frac{5}{8}\mu \log n$. Thus, using the Chernoff bound, with probability

$$1 - n(4\mu n \log n) e^{-\frac{1}{3}\frac{9}{16}(\frac{1}{9})^2 \mu \log n}, \tag{4}$$

---

[2] We say that distribution F is stochastically lower- (upper-) bounded by distribution G if for any x, $F(x) \geq G(x)$ (respectively, $F(x) \leq G(x)$).

either $X_\tau \le 6\alpha \log n$ for some τ in that interval or no sender had a success rate of at least $\frac{5}{8}$ in a sequence of $\mu \log n$ attempts in that interval.

Combining the bounds in (2) and (4) we get that for each interval of $4\mu n \log n$ steps, with probability at least

$$1 - e^{-\mu \log n} - n(4\mu n \log n) e^{-\frac{1}{3}\frac{9}{16}(\frac{1}{9})^2 \mu \log n} \ge 1 - n^{-12}$$

(for sufficiently large constant μ), either $X_\tau \le 6\alpha \log n$ at some step of the interval or at least one sender switches to the reset state and moves all its packets from its buffer to its queue. Thus, with probability $1 - n^{-11}$, at some step $\tau_1 < t + 4\mu n^2 \log n$, $X_{\tau_1} \le 6\alpha \log n$. $\square$

CLAIM 2. *Suppose that there is a positive integer L such that if $X_\tau \in [L, L + 2\alpha \log n]$ then the probability that a packet is delivered from a buffer at step τ is at least p for a constant $p > \lambda$. If $X_{\tau_1} \le L$, then*

$$\mathbf{Pr}[X_\tau \ge L + \alpha \log n \textit{ for some step } \tau \in [\tau_1, \tau_1 + W]] \le W n^{-13}$$

for a sufficiently large constant α.

Proof. Since at most one message can be delivered from all buffers in each step, X_τ can decrease by at most one in each step. Thus, it is sufficient to prove that there is no time interval of length $\alpha \log n$ during which X_τ remains above L.

Consider an interval of $\alpha \log n$ steps, starting at step $\tau_2 \in [\tau_1, \tau_1 + W]$. Assume that $X_{\tau_2 - 1} \le L$. Let $\epsilon = (p - \lambda)/3\lambda$. The probability that more than $(1+\epsilon)\lambda\alpha \log n < 2\alpha \log n$ packets are placed in buffers in the interval is at most $n^{-13}/2$.

Let z_τ be a 0–1 random variable defined as follows: $z_\tau = 1$ if $X_\tau \le L$ or a packet is delivered from a buffer in step τ; else $z_\tau = 0$. Clearly, as long as $X_\tau \le L + 2\alpha \log n$, $\mathbf{Pr}[z_\tau = 1 \mid z_1, \ldots, z_{\tau-1}] \ge p$. Let $\delta = (p - \lambda)/3p$.

$$\mathbf{Pr}\left[\sum_{k=\tau_2}^{\tau_2 + \alpha \log n} z_k < (1 - \delta) p \alpha \log n\right] \le e^{-\frac{1}{3}\delta^2 p \alpha \log n} \le n^{-13}/2.$$

Since $(1 + \epsilon)\lambda\alpha \log n < (1 - \delta)p\alpha \log n$, the probability that $X_\tau > L$ for a given $\tau \in [\tau_2, \tau_2 + \alpha \log n]$ is at most n^{-13}, for any $p > \lambda$ and a sufficiently large constant α. $\square$

CLAIM 3. *If $X_{\tau_1} \le 6\alpha \log n$ for some $\tau_1 \in [t - T, t - \gamma \log n]$, then with probability at least $1 - n^{-11}$, there is a step $\tau_2 \in [\tau_1, t]$ such that $X_{\tau_2} \le \alpha \log n$.*

Proof. As long as $X_\tau \in [\alpha \log n, 7\alpha \log n]$ the probability of a successful transmission from some buffer at time τ is at least p, where

$$p = \binom{X_\tau}{1} \frac{1}{8\alpha \log n}\left(1 - \frac{1}{8\alpha \log n}\right)^{X_\tau - 1} - \frac{n}{n^2} \ge \frac{X_\tau}{8\alpha \log n} - (\frac{X_\tau}{8\alpha \log n})^2 - \frac{n}{n^2} \ge \frac{1}{10}.$$

Let $\mathcal{E}_1$ denote the event: "$X_{\tau_1} \le 6\alpha \log n$, and there is a $\tau \in [\tau_1, \tau_1 + \gamma \log n]$ such that $X_\tau > 7\alpha \log n$."

By Claim 2, $\mathbf{Pr}[\mathcal{E}_1] \le n^{-13}\gamma \log n$ for any $\lambda < p$ and a sufficiently large constant α.

The expected number of new packets arriving in the time interval $[\tau_1, \tau_1 + \gamma \log n]$ is $\lambda\gamma \log n$. Let $\delta = (p - \lambda)/3\lambda$ and let $\mathcal{E}_2$ denote the event "More than $\lambda\gamma(1+\delta) \log n$ packets arrived in the interval $[\tau_1, \tau_1 + \gamma \log n]$." Then $\mathbf{Pr}[\mathcal{E}_2] \le e^{-\delta^2 \lambda\gamma \log n/3} \le n^{-11}$ for a sufficiently large γ.

Let $\epsilon = (p - \lambda)/3p$ and define the event $\mathcal{E}_3$: “either $X_\tau \leq \alpha \log n$ for some $\tau \in [t - \gamma \log n, t]$ or at least $(1-\epsilon)p\gamma \log n$ packets were delivered from the buffer at that interval.” Then $\mathbf{Pr}[\mathcal{E}_3 \mid \bar{\mathcal{E}_1}] \geq 1 - e^{-\epsilon^2 p\gamma \log n/3}$.

Fix γ such that $\gamma(p-\lambda)/3 > 6\alpha$. Then the probability that there is no $\tau_2 \in [\tau_1, \tau_1 + \gamma \log n]$ such that $X_{\tau_2} < \alpha \log n$ is at most

$$n^{-13}\gamma \log n + e^{-\lambda\gamma\delta^2 \log n/3} + e^{-p\gamma\epsilon^2 \log n/3} \leq n^{-11}$$

for $\lambda < p$ and $\gamma \geq \max[\frac{18\alpha}{p-\lambda}, \frac{39}{\lambda\delta^2}, \frac{39}{p\epsilon^2}]$. □

CLAIM 4. *If $X_{\tau_1} \leq \alpha \log n$ then the probability that there exists $\tau \in [\tau_1, \tau_1 + T]$ such that $X_\tau > 2\alpha \log n$ is at most n^{-11}.*

Proof. If $X_\tau \in [\alpha \log n, 4\alpha \log n]$ then the probability of a successful transmission from a buffer at time τ is at least

$$p = \binom{X_t}{1}\frac{1}{8\alpha \log n}\left(1 - \frac{1}{8\alpha \log n}\right)^{X_t - 1} - \frac{n}{n^2}$$

$$\geq \frac{X_t}{8\alpha \log n} - \left(\frac{X_t}{8\alpha \log n}\right)^2 - \frac{n}{n^2} \geq \frac{1}{10}.$$

Using Claim 2, the probability of having $2\alpha \log n$ packets in buffers in the interval is at most Tn^{-13} for any $\lambda < p$ and a sufficiently large constant α. □

To conclude the proof of Lemma 1 we combine the error probabilities of the three lemmas above to show that for any value x,

$$\mathbf{Pr}[X_t > 2\alpha \log n \mid X_{t-T} = x] \leq 3n^{-11} \leq n^{-10}. \qquad \square$$

Comment. The statement of Lemma 1 implies that in the steady state, almost always $X_t \leq 2\alpha \log n$. This observation is not used directly in the proof but gives a good intuition for the long-term behavior of the protocol.

We turn to the analysis of the queues. We focus on the performance of the protocol after the first $T_0 = 4\mu n^2 \log n + \gamma \log n + 4\mu n \log n$ steps. We assume that at time 0 all queues are empty (this assumption affects only the length of the initial segment T_0).

CLAIM 5. *Let $t \geq T_0$. The probability that any sender switches to the reset state at step t is at most n^{-8}.*

Proof. If $X_\tau \leq 2\alpha \log n$, then the probability of a success in a transmission attempt at step τ is at least $p = 1 - (2\alpha \log n)/(8\alpha \log n) - (n/n^2) \geq 3/4 - 1/n$.

For a given sender s and given steps $\tau \geq \tau' \geq t - 4\mu n \log n$, let $y_\tau^{\tau',s}$ be a random variable defined as follows: $y_\tau^{\tau',s} = 1$ if sender s successfully transmitted a packet at time τ, $X_\tau \leq 2\alpha \log n$, and s had less than $\mu \log n$ transition attempts in the interval $[\tau', \tau]$; else $y_\tau^{\tau',s} = 1$. Let $\mathcal{H}_\tau$ describe the state of the system at all times before step τ. Clearly

$$\mathbf{Pr}[Y_\tau^{\tau',s} \mid \mathcal{H}_\tau] \leq \frac{1}{4} + \frac{1}{n},$$

and $\sum_{\tau=\tau'}^{\tau_0+4\mu n \log n - 1} y_\tau^{\tau',s}$ is stochastically upper-bounded by a binomial distribution with expectation $(\frac{1}{4} + \frac{1}{n})\mu \log n$. Thus, by the Chernoff bound with probability

$$1 - n(4\mu n \log n)e^{-\frac{1}{3}(\frac{1}{4}+\frac{1}{n})(\frac{1}{2})^2\mu \log n}, \tag{5}$$

either $X_\tau > 2\alpha \log n$ for some τ in that interval or no sender had a success rate less than $\frac{5}{8}$ in a sequence of $\mu \log n$ attempts in that interval. By Lemma 1, $\mathbf{Pr}[X_\tau > 2\alpha \log n]$ at any time $\tau \in [t-4\mu n \log n, t]$ is at most $4\mu n^{-9}(\log n+2)$, which together with the above bound proves the lemma. ☐

LEMMA 2. *Consider a packet that was generated at time $t \geq T_0$.*

1. *The expected number of steps a given packet spends in a buffer is $O(\log n)$.*

2. *The probability that a given packet is delivered from the buffer is at least $1 - n^{-5}$.*

Proof. Assume that the packet was generated at sender s.

Consider an interval of n steps throughout which (1) s is never in the reset state, and (2) $X_\tau \leq 2\alpha \log n$. Let $p = 1 - (2\alpha \log n)/(8\alpha \log n) - n/n^2$ as in the proof of Claim 5. If the packet was in the buffer of sender s at the start of the interval, the probability that it is not delivered until the end of the interval is at most $(1 - p/8\alpha \log n)^n$. During the interval the expected number of steps between two attempts to transmit the packet is $8\alpha \log n$, and the expected number of attempts until the packet is delivered is $1/p$. Thus, if the packet is delivered during the interval, by Wald's identity [2] the expected delay of the packet from the start of that interval is at most $(8\alpha \log n)/p$.

The probability that the interval $[t, t+n]$ satisfies conditions (1) and (2) is (by Lemma 1 and Claim 5) at least $1-(4\mu n^2 \log n+n)n^{-8}-n^{-9}$, and with this probability, any subsequent interval of $4\mu n^2 \log +n$ steps has a segment of n steps that satisfies the above conditions. Thus, the expected number of steps a packet spends in the buffer is at most

$$\frac{8\alpha \log n}{p} + \sum_{k\geq 1} k(n + 4\mu n^2 \log n) \left((4\mu n^2 \log n + n)n^{-8} + n^{-9} + \left(1 - \frac{p}{8\alpha \log n}\right)^n \right)^k$$

$$= \frac{8\alpha \log n}{p} + o(1).$$

(The above estimate ignores the possibility that a packet is moved to the queue before it is delivered from the buffer. This can only decrease the expected number of steps a packet spends in the buffer.)

The probability that the packet is delivered from the buffer is bounded from below by the probability that the first n steps following the creation of the packet satisfy conditions (1) and (2) and the packet is delivered in that time. This probability is at least

$$1 - (4\mu n^2 \log n + n)n^{-8} - n^{-9} - \left(1 - \frac{p}{8\alpha \log n}\right)^n \geq 1 - n^{-5}. \qquad ☐$$

LEMMA 3. *Consider a packet that was generated at time $t > T_0$. Given that the packet enters the queue, the expected length of time it spends in the queue is $O(n^4 \log n)$.*

Proof. Assume that the packet was generated at sender s. When a sender switches to the reset state it moves all the (up to $2\alpha \log n + 1$) packets in its buffer to its queue. In addition, all the (up to $4\mu n^2 \log n + \gamma \log n$) packets it receives in the next $4\mu n^2 \log n + \gamma \log n$ steps are placed in the queue. Let $D = 4\mu n^2 \log n + \gamma \log n + 2\alpha \log n + 1$.

When a queue is not empty and the total number of packets in all buffers in the system is at most $2\alpha \log n$, a sender succeeds in transmitting a packet from its queue in each step with probability at least

$$p = \frac{1}{n^2}\left(1 - \frac{1}{8\alpha \log n}\right)^{2\alpha \log n}\left(1 - \frac{1}{n^2}\right)^{n} \geq \frac{1}{2n^2}.$$

Let τ_1 be the first step that X_t exceeds $2\alpha \log n$ (i.e., $X_{\tau_1 - 1} \leq 2\alpha \log n$ and $X_{\tau_1} > 2\alpha \log n$). Let τ_i, $i > 1$, be the first step that X_t exceeds $2\alpha \log n$ after step $\tau_{i-1} + T$ ($T = 4\mu n^2 \log n + \gamma \log n$ as defined in Lemma 1). We say that an interval $[\tau_i, \tau_{i+1}]$ is *good* (with respect to sender s) if either the queue of s was empty at some step in that interval or at least $2D$ packets were delivered from the queue. Note that no more than D packets can enter the queue in each interval, since a sender cannot switch twice to a reset state in the same interval.

By Lemma 1 the probability that an interval does not have a segment of $L = 4n^2 D$ steps in which $X_t \leq 2\alpha \log n$, conditioning on all events before that segment is bounded by Ln^{-10}. The probability that in L steps in which $X_t \leq 2\alpha \log n$, fewer than $2D$ packets are transmitted, is bounded by $e^{-\frac{1}{3}\frac{1}{2n^2}L} \leq e^{-\frac{2D}{3}}$. Thus, in a given interval the probability that either $Z_t = 0$ at some point in the interval or at least $2D$ packets are delivered from the queue is at least $1 - Ln^{-10} - e^{-\frac{2D}{3}} \geq 1 - \frac{1}{n^5}$, and for subsequent intervals these events are independent.

The probability that there are iD packets in the queue at time t is bounded by the probability that there are fewer than $(k-i)/2$ *good* intervals in the k most recent intervals, for some $k \geq i$. This probability is bounded by $\sum_{k \geq i} \binom{k}{(k+i)/2}(\frac{1}{n^5})^{(k+i)/2} \leq 2(\frac{e}{n^5})^i$, and the expected number of packets in the queue at time t is bounded by $\sum_{i \geq 0} iD(\frac{e}{n^5})^i = O(1)$. Our packet, however, arrives with another $D = O(n^2 \log n)$ packets, and the expected number of steps to send a packet from the head of the queue is $O(n^2)$. Thus, the expected time until all these packets are transmitted is $O(n^4 \log n)$. □

By Claim 2 the expected time a packet spends in the buffer in $O(\log n)$. With probability $O(n^{-5})$ the packet is transferred to the queue, where its expected delay is $O(n^4 \log n)$. Thus, the expected delay of a given packet is $O(\log n)$. □

3. The protocol for n senders and m receivers.

In our protocol each sender runs m "one-channel protocols" simultaneously. Each sender keeps a counter of transmission failures for each of the receivers. When this count is too large the sender switches to the reset state only with respect to packets bound for that receiver.

The "one-channel protocol" requires that if there are X packets in the buffer of sender s, the sender tries to transmit a packet from the buffer with probability $X/8\alpha \log n$. To accommodate the m protocols simultaneously in one sender we modify step 2 in the "one-channel protocol" as follows:

2. Let $Y^i(s)$ denote the number of packets with destination i in the buffer of sender s. Let $Y(s) = \sum_{i=1}^{m} Y^i(s)$.
If $Y(s) > 2\alpha \log n$ **then** move all packets to the queue,
else with probability $Y(s)/8\alpha \log n$ transmit a random packet from the buffer.

It remains to show that the probability that a sender has more than $2\alpha \log n$ packets in the buffer is small. Thus, moving them to the queue does not significantly change the performance of the protocol.

Lemma 4. *Let $Y_t(s)$ denote the total number of packets in the buffer of sender s at time t. In the stationary distribution,*

$$\mathbf{Pr}[\exists s \text{ such that } Y_t(s) > 2\alpha \log n] \leq n^{-7}.$$

Proof. As in the proof of Lemma 1 we show that there exists a constant γ such that if $T = 4\mu n^2 \log n + \gamma \log n$, then for any value y of $Y_{t-T}(s)$,

$$\mathbf{Pr}[Y_t(s) > 2\alpha \log n \mid Y_{t-T}(s)] = y \leq n^{-8}.$$

Let X^i_τ denote the total number of packets with destination i in the buffers of all senders. Let $\mathcal{E}$ denote the event: "For any $1 \leq i \leq m$ and $t - \beta \log n \leq \tau \leq t$, $X^i_\tau \leq 2\alpha \log n$."

By Lemma 1, regardless of the state of the system at time $t - T$, $\mathbf{Pr}[\mathcal{E}] \geq 1 - n^{-9}$ (since there are no more than n protocols).

The total injection rate of sender s is at most λ. Conditioning on the event $\mathcal{E}$, the probability that s successfully delivers a message at time τ, $\tau \in [t - \beta \log n, t]$ when $Y_\tau(s) > \alpha \log n$ is at least $p = \frac{\alpha \log n}{8\alpha \log n}(1 - 1/8\alpha \log n)^{2\alpha \log n}(1 - 1/n^2)^n \geq 1/20$. As in the proof of Claim 3 we show that with probability $1 - n^{-9}$ there is a step $\tau \in [t - \beta \log n, t]$, such that $Y_t(s) < \alpha \log n$, and using Claim 2 we show that once the number of packets is below $\alpha \log n$ the probability that this number reaches $2\alpha \log n$ before time T is at most n^{-8}. □

When a sender has $2\alpha \log n$ packets in its buffer it transfers them to its queue. Thus, the expected additional contribution to the queue from the modified protocol is $O(n^{-8} \log n)$ per step. This contribution does not change the delivery time of packets in the queue or the overall performance of the protocol significantly.

Theorem 2. *There is a constant $\lambda_0 > 0$ such that the above protocol is stable for any $0 \leq \lambda \leq \lambda_0$ and the delay is $O(\log n)$.*

4. The lower bound.

We state and prove the lower bound for the case $m = 1$; a similar bound holds for the general case. We consider n senders all running the same protocol. The protocol running at a sender does not change state unless (1) it attempts a transmission or (2) it generates a packet. We further assume that a sender is in a unique idle state when its queue is empty. (Thus, a sender does not keep information on the state of the channel when it does not have packets to transmit.) For any such protocol governed by a (probabilistic) automaton we show that there is a $\lambda_0 < 1$ such that for any injection rate $\lambda > \lambda_0$ the delay is $\Omega(n)$. Note that the class of protocols includes all variants of backoff protocols.

Theorem 3. *If all senders follow protocols from the above class, there is a fixed constant $\lambda_0 < 1$ such that if $\lambda > \lambda_0$, the expected delay is $\Omega(n)$.*

Proof. Let $\lambda_i = \lambda/n$ for all i. Given a system of senders, we say that $X_t = j$ if the total number of packets pending transmission at all senders at step t is j. If the expected delay is bounded then the system has a steady state distribution. Consider the steady state distribution of X_t. We distinguish between two cases.

Case 1. In the steady state distribution, $\mathbf{Pr}[X_t] \geq n/2 \geq 1/2$. Since the system can transmit no more than one packet per step the expected delay in this case in clearly $\Omega(n)$.

Case 2. In the steady state distribution $\mathbf{Pr}[X_t] < n/2 > 1/2$. □

Lemma 5. *If $X_t < n/2$, the probability of a collision at step $t + n/3$ is at least a constant q.*

Proof. If $X_t < n/2$, we have at least $n/2$ senders having no packets pending transmission at time t (call these *empty senders*). A protocol in our class defines a probability distribution p_i on the nonnegative integers i such that when an empty sender generates a packet at step t, it transmits it at step $t+i$ with probability p_i unless it changes its behavior (due to the generation of a packet) in the interval $[t, t+i]$.

An empty sender that generates a packet at time τ does not change the probabilities p_i in the interval $(\tau, \tau + n/3]$ unless it generates another packet in the interval. The probability that it generates such a packet in the interval $(\tau, \tau+n/3]$ is at most $(1-\lambda/n)^{n/3}$, which is at most a constant α. Thus the empty sender does transmit the packet at time $\tau + i$ with probability at least $(1-\alpha)p_i$. Clearly, $\sum_{i=0}^{n/3} p_i$ must be at least a constant β (for otherwise the delay is already at least $\alpha(1-\beta)n/3$). If at least $n/2$ senders are empty at time t, the probability that any one of them transmits at time $t + n/3$ is thus at least

$$(1-\alpha)\sum_{k=0}^{n/3} \frac{\lambda}{n} p_k \geq (1-\alpha)\beta\lambda/n.$$

Then, the probability of a collision at time $t + n/3$ is at least a constant q. ∎

Thus, if in the steady state distribution, $\mathbf{Pr}[X_t] < n/2 > 1/2$, then the probability of a collision in a given step in the steady state is at least a constant $q/2$. Since the expected number of new packets arriving in each step is λ and the system can transmit no more than one packet per step, the system cannot be stable if $\lambda > 1 - q/2$.

Acknowledgment. We thank Aravind Srinivasan for helpful comments.

REFERENCES

[1] D. Aldous, *Ultimate instability of exponential back-off protocol for acknowledgment based transmission control of random access communication channels*, IEEE Trans. Inform. Theory, 33 (1987), pp. 219–223.

[2] W. Feller, *An Introduction to Probability Theory and Its Applications, Vol.* II, John Wiley, New York, 1968.

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

[4] A.G. Greenberg, P. Flajolet, and R.E. Ladner, *Estimating the multiplicities of conflicts to speed their resolution in multiple access channels*, J. Assoc. Comput. Mach., 34 (1987), pp. 289–325.

[5] L.A. Goldberg and P.D. MacKenzie, *Analysis of practical backoff protocols for contention resolution with multiple servers*, in Proceedings of the 7th Annual Symp. on Discrete Algorithms, SIAM, Philadelphia, 1996, pp. 554–563.

[6] J. Håstad, T. Leighton, and B. Rogoff, *Analysis of backoff protocols for multiple access channels*, in Proceedings of the 19th ACM Symp. on Theory of Computing, 1987, pp. 241–253.

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

[8] P.D. MacKenzie, C.G. Plaxton, and R. Rajaraman, *On contention resolution protocols and associated probabilistic phenomena*, in Proceedings of the 26th ACM Symp. on Theory of Computing, 1994, pp. 153–162.

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

[10] R. Motwani and P. Raghavan, *Randomized Algorithms*, Cambridge University Press, New York, 1995.

[11] M. Paterson and A. Srinivasan, *Contension resolution with bounded delays*, in Proceedings of the 36th Annual ACM Symp. on Foundations of Computer Science, 1995, pp. 104–113.