

UC Berkeley

UC Berkeley Previously Published Works

Title

Ultimate instability of exponential back-off protocol for acknowledgment-based transmission control of random access communication channels

Permalink

<https://escholarship.org/uc/item/09b8w9pn>

Journal

IEEE Transactions on Information Theory, 33(2)

ISSN

0018-9448

Author

Aldous, D

Publication Date

1987-03-01

DOI

10.1109/tit.1987.1057295

Peer reviewed

Ultimate Instability of Exponential Back-Off Protocol for Acknowledgment-Based Transmission Control of Random Access Communication Channels

DAVID J. ALDOUS

Abstract—When several users simultaneously transmit over a shared communication channel, the messages are lost and must be retransmitted later. Various protocols specifying when to retransmit have been proposed and studied in recent years. One protocol is “binary exponential back-off,” used in the local area network *Ethernet*. A mathematical model with several idealizations (discrete time slots, infinite users, no deletions) is shown to be unstable in that the asymptotic rate of successful transmissions is zero, however small the arrival rate.

I. INTRODUCTION

MULTIACCESS communication channels have been the object of intense study in recent years, motivated by problems of communications between computers. A special journal issue [6] provides an overview of the field; Kelly [7, sec. 4] gives a concise account of the specific problem treated here. Consider N geographically separated “users” who can communicate with each other via one shared channel to which all users are constantly listening. When one user has a message it wishes to send to another, it can transmit the message over the channel, and the message will be successfully received provided no other user is simultaneously transmitting, in which case none of the conflicting transmissions are successfully received. In designing such a system one might seek to avoid such conflicts by having a centralized controller who schedules transmissions; the difficulty is that users would then have to tell the controller when they have messages to send, and they have to use the channel to send this information! Instead, it is easier to implement systems where each user acts autonomously: if a transmission is unsuccessful, the user waits a random time (determined by some strategy) and then retransmits the message, continuing until a transmission is successful.

A standard simplified model is as follows. Divide time into slots $(t, t + 1)$, $t = 0, 1, 2, \dots$, of unit length, and

suppose all messages take unit time. Each user originates a new message in each time slot $(t - 1, t)$ with probability p (independently over users and time slots); such a new message is transmitted as soon as possible, that is, in slot $(t, t + 1)$. The strategy for retransmitting is described by a vector $(h_i; 1 \leq i < K)$, where $0 < h_i < 1$ and $K \leq \infty$. In slot $(t, t + 1)$ each message which has had i previous unsuccessful transmissions is selected for transmission with probability h_i (independently over messages and time slots). After K unsuccessful transmissions a message is deleted and declared lost. This kind of control policy is called “acknowledgment based” (or collision detect) because the only information a user requires is the acknowledgment of successful receipt of its own messages. More sophisticated policies, where each user monitors the channel continuously and notes whether 0, 1, or more than one transmission is attempted in each slot, have been studied [6] but will not be treated here.

To avoid the possibility that one user may have more than one message to transmit at one time, it is convenient to pass to the “infinite-users” model in which we let $N \rightarrow \infty$, $p \rightarrow 0$ and $Np \rightarrow \nu > 0$. The state of the system at time t is described by a vector $x = (x_i; 1 \leq i < K)$ where x_i is the number of messages which have been unsuccessfully transmitted exactly i times. A random subset of these messages (selected as described earlier using (h_i)) together with a Poisson (ν) number of new messages (those originating during $(t - 1, t)$) are transmitted during $(t, t + 1)$; if exactly one transmission is made, then it is successfully received, otherwise, all transmissions are unsuccessful. This describes a countable state-space Markov chain $X(t) = (X_i(t); 1 \leq i < K)$ whose transition probabilities depend only on ν and (h_i) .

Studies of this and closely related models, with particular attention to the special cases $h_i \equiv h$ (the ALOHA policy) and $h_i = 2^{-i}$ (binary exponential back-off, or the ETHERNET policy) have been given in [2]–[9] and the papers referenced therein. Let $N(t)$ be the number of successful transmissions made during time $[0, t]$. For $K < \infty$ it is easy to show that X is positive recurrent, though no useful form of the stationary distribution is known. It

Manuscript received March 1986; revised June 1986. This work was supported in part by the National Science Foundation under Contract MCS80-02698.

The author is with the Statistics Department, University of California, Berkeley CA 94720, USA.

IEEE Log Number 8610918.

follows that the asymptotic rate of successful transmissions

$$\rho = \text{a.s.} \lim_{t \rightarrow \infty} N(t)/t$$

exists, and $0 < \rho < \nu$. Therefore, a proportion $1 - \rho/\nu$ of messages are lost. From now on, consider the case where $K = \infty$, that is, messages are never deleted. It is natural to hope that, for some choice of (h_i) , the process X will still be positive recurrent and $\rho = \nu$. However, some negative results are known. Kelly [7] and Kelly and MacPhee [8] give a formula for a critical value ν_c (depending on (h_i)) such that

$$\begin{aligned} \lim_{t \rightarrow \infty} N(t) &= \infty \text{ a.s.,} & \nu < \nu_c \\ \lim_{t \rightarrow \infty} N(t) &< \infty \text{ a.s.,} & \nu > \nu_c. \end{aligned} \quad (1.1)$$

For the ALOHA scheme $\nu_c = 0$; for binary exponential back-off $\nu_c = \log 2$. Obviously, this implies $\rho = 0$ for $\nu > \nu_c$. However, this result and related results of Fayolle [3] and Rosenkrantz [9] leave open the behavior of ρ for small ν in the binary exponential back-off case.

Theorem 1: For $K = \infty$ and any $\nu > 0$ the binary exponential back-off policy is unstable in the sense that the chain $X(t)$ is transient and $N(t)/t \rightarrow 0$ a.s.

The proof is given in Section II with technical lemmas deferred to Section III. The argument is rather more delicate than previous cases, though some of the ingredients are the same, e.g., comparisons with the externally jammed channel. We end with a series of remarks.

a) Without checking the details, I believe the argument could be modified to show instability for every (h_i) : the essential change would be in the definition of (t_i) to follow in (2.5). A harder problem is to show that every acknowledgment-based policy is unstable since one could invent policies much more complicated than those considered here.

b) A slight tightening of our argument shows

$$N(t) = o(t^a) \text{ a.s.,} \quad \text{all } a > 1 - \nu/\log 2. \quad (1.2)$$

Note that this is consistent with (1.1). These results suggest that for $\nu < \log 2$ the asymptotic growth rate of $N(t)$ is on the order of $t^{1-\nu/\log 2}$. However, asymptotics are rather misleading. Simulations and heuristic arguments [2], [4], [5] suggest that, for fairly general (h_i) , if ν is rather smaller than $1/e$, then the process $X(t)$ quickly reaches a quasistationary distribution which persists for a long, but finite, time T_s before instability sets in. Formalizing this and estimating ET_s is another challenging problem.

c) The significant idealizations made in our model are 1) time is divided into discrete slots 2) infinite users, and 3) no deletions of unsuccessfully transmitted messages. Discussions of models without one of these simplifications are given in [8], [4], and [5], respectively.

II. CONSTRUCTION AND PROOF

We start by giving a "balls in boxes" description of our process. Imagine boxes $1, 2, 3, \dots$ in a line left to right, and another box 0 containing an infinite supply of balls.

The state of the process is a vector $y = (y_i; i \geq 1)$ indicating y_i balls in box i . Given state y at time t , the state at time $t + 1$ is determined as follows. Each ball in boxes $i \geq 1$ is either moved one box to the right (with chance 2^{-i}) or remains in box i , independently for different balls. Further, a Poisson (ν) number of balls are moved from box 0 to box 1. Here $\nu > 0$ is a fixed parameter. This specifies a Markov process $Y(t) = (Y_i(t); i \geq 1), t \geq 0$. The fact that balls progress independently makes it easy to verify that $Y(t)$ has a stationary distribution $Y = (Y_i)$ given by

$$(Y_i; i \geq 1) \text{ are independent,} \quad Y_i \stackrel{\mathcal{D}}{=} \text{Poisson}(\nu 2^i). \quad (2.1)$$

Give $Y(0)$ this stationary distribution. We now introduce a coloring scheme. At time 0 let the balls in boxes $i \geq 1$ be colored red and the balls in box 0 be colored white. Therefore, balls keep their color, except for the proviso: if the set of balls moved at one time contains exactly one white ball, then that white ball is colored red. Let $X_i(t)$ be the number of white balls in box i at time t , and let $X(t) = (X_i(t); i \geq 1)$. Therefore, $X(t)$ is a nonstationary Markov process, and a moment's thought shows it is precisely the same as the binary exponential back-off process of Section I with the white balls in box i corresponding to messages which have been unsuccessfully transmitted i times and the "recolorings" corresponding to successful transmissions. Moreover, $Y(t)$ evolves as the externally jammed process where no transmissions are successful.

Note that the inequality $X_i(t) \leq Y_i(t)$ implies $EX_i(t) \leq EY_i(t) = \nu 2^i$, yielding a simpler proof of bounds obtained in [9]. The basic idea of our proof is as follows. For the externally jammed channel Y we have $EY_i(t) = \nu 2^i$. For the real channel X , even if it were stable, occasional times s would occur when "by chance" $X_i(s) \geq \nu 2^i$ for an arbitrarily long block of i 's. We shall show (2.5) that when this happens, a positive chance exists that the channel becomes more and more jammed.

We now start the mathematical analysis, deferring proofs of lemmas. Let

$$f(x) = \sum_{i=1}^{\infty} x_i 2^{-i}. \quad (2.2)$$

Lemma 1: $P(\text{some ball recolored on } (t+1)\text{th move} | X(t) = x) \leq 2 \exp(-(1/2)f(x))$.

We shall prove

$$f(X(t)) \rightarrow \infty \quad \text{a.s. as } t \rightarrow \infty. \quad (2.3)$$

Then $X(t)$ is transient, and using Lemma 1, the asymptotic rate of recolorings is a.s. zero, proving Theorem 1.

Let L_0 be an integer, sufficiently large to satisfy various constraints we shall specify later. For $s \geq 2^{L_0+5}$ let A_s be the event

$$X_i(s) \geq \nu 2^i, \quad \text{for all } 1 \leq i \leq L_0. \quad (2.4)$$

Lemma 2: $P(A_s \text{ happens infinitely often}) = 1$.

Let \mathcal{F}_s be the σ -field generated by $(X(u), Y(u); u \leq s)$. We shall prove that for each s

$$P(f(X(t)) \rightarrow \infty \text{ as } t \rightarrow \infty | \mathcal{F}_s) \geq \frac{1}{2} \text{ on } A_s. \quad (2.5)$$

Writing B for the event $f(X(t)) \rightarrow \infty$ as $t \rightarrow \infty$, Lemma 2 and (2.5) imply

$$\limsup_{s \rightarrow \infty} P(B | \mathcal{F}_s) \geq \frac{1}{2} \text{ a.s.}$$

However, the martingale convergence theorem [1, p. 93] says

$$P(B | \mathcal{F}_s) = E(1_B | \mathcal{F}_s) \rightarrow 1_B, \quad \text{a.s. as } s \rightarrow \infty.$$

Hence $1_B \geq 1/2$ a.s.; in other words, $P(B) = 1$. This gives (2.3) and thence Theorem 1.

To prove (2.5), fix s and define inductively

$$\begin{aligned} t_0 &= s \\ t_i - t_{i-1} &= 2^{L_0+6+i}. \end{aligned}$$

For $t \geq s$ define

$$\begin{aligned} L(t) &= L_0 + i \text{ on } t_i \leq t < t_{i+1} \\ C_i &= \left\{ f(X(t)) \geq \frac{1}{2} \nu L(t) \right\} \\ B_i &= \bigcap_{s \leq u \leq t} C_u. \end{aligned}$$

We shall prove that

$$\sum_{t \geq s+1} P\left(f(X(t)) < \frac{1}{2} \nu L(t), B_{i-1} | \mathcal{F}_s\right) \leq \frac{1}{2} \text{ on } A_s. \quad (2.6)$$

Now this sum is

$$\sum_{t \geq s+1} \{P(B_{i-1} | \mathcal{F}_s) - P(B_i | \mathcal{F}_s)\},$$

and $P(B_s | \mathcal{F}_s) = 1$ on A_s by construction, so (2.6) implies

$$P\left(\bigcap_{u \geq s} C_u | \mathcal{F}_s\right) = \lim_{t \rightarrow \infty} P(B_t | \mathcal{F}_s) \geq \frac{1}{2} \text{ on } A_s,$$

which implies (2.5).

So far we have argued “backwards” and shown that it suffices to prove (2.6). We now start arguing “forwards.” For $t \geq t_0$ let $Z_i(t)$ be the number of balls in box i at time t which were in box 0 at time $t - 2^{4+i}$. Let $D_i(t)$ be the total number of recolorings during the interval $[t - 2^{4+i}, t]$. Then

$$X_i(t) \geq Z_i(t) - D_i(t) \quad (2.7)$$

because any ball in box i at time t which left box 0 after time $t - 2^{4+i}$ and has not been recolored must be white. Write

$$\alpha(L) = \exp(-\nu 2^{L/10}/200). \quad (2.8)$$

Lemma 3: If $t - t_0 \geq 2^{4+i}$ and $i \geq L(t)/10$, then

$$P\left(D_i(t) \geq \frac{1}{10} \nu 2^i, B_{i-1} | \mathcal{F}_{t_0}\right) \leq \alpha(L(t)).$$

Lemma 4: If $t - t_0 \geq 2^{4+i}$ and $i \geq L(t)/10$, then

$$P\left(Z_i(t) \leq \frac{7}{9} \nu 2^i\right) \leq \alpha(L(t)).$$

Noting that $Z_i(t)$ is independent of \mathcal{F}_{t_0} if $t - t_0 \geq 2^{4+i}$, we can combine (2.7) and Lemmas 3 and 4 to get

$$P\left(X_i(t) \leq \frac{61}{90} \nu 2^i, B_{i-1} | \mathcal{F}_{t_0}\right) \leq 2\alpha(L(t)),$$

$$\text{provided } t - t_0 \geq 2^{4+i}, \quad i \geq L(t)/10. \quad (2.9)$$

To cover the case where $t - t_0$ is small, a separate argument exists for the following lemma.

Lemma 5: If $0 \leq t - t_0 \leq 2^{i-5}$ and $L_0/10 \leq i \leq L_0$, then

$$P\left(X_i(t) \leq \frac{9}{10} \nu 2^i | \mathcal{F}_{t_0}\right) \leq \alpha(L_0) \text{ on } A_s.$$

Given $t \geq t_0$, consider for how many i 's in the range $L(t)/10 \leq i \leq L(t)$ the inequality (2.9) is true. If $t \geq t_1 = t_0 + 2^{L_0+6}$, then the condition $t - t_0 \geq 2^{4+i}$ holds for all $i \leq L(t)$ by construction of $L(t)$. If $t < t_1$, then $L(t) = L_0$, and at most nine values of i exist in the desired range for which the conditions of neither (2.9) nor Lemma 5 are satisfied. Since the conclusion of Lemma 5 implies the weaker assertion (2.9) on A_s , we have proved the following: given $t \geq t_0$, the inequality

$$P\left(X_i(t) \leq \frac{61}{90} \nu 2^i, B_{i-1} | \mathcal{F}_{t_0}\right) \leq 2\alpha(L(t)) \text{ on } A_s \quad (2.10)$$

holds for at least $(9/10)L(t) - 9$ values of i .

Now $f(X(t)) = \sum_{i=1}^{\infty} X_i(t) 2^{-i}$, and by taking L_0 sufficiently large,

$$\frac{61}{90} \left(\frac{9}{10} L - 9\right) \geq \frac{1}{2} L, \quad L \geq L_0.$$

Therefore, (2.10) implies that for $t > t_0 = s$,

$$\begin{aligned} P\left(f(X(t)) \leq \frac{1}{2} \nu L(t), B_{i-1} | \mathcal{F}_{t_0}\right) \\ \leq 2L(t)\alpha(L(t)) \text{ on } A_s. \end{aligned}$$

Thus to prove (2.6), it suffices to show that for sufficiently large L_0

$$\sum_{t \geq t_0} L(t)\alpha(L(t)) < \frac{1}{4},$$

and this is just calculus.

III. TECHNICAL LEMMAS

Recall the elementary large deviation bounds for an arbitrary random variable Z :

$$P(Z \geq z) \leq \inf_{\theta > 0} e^{-\theta z} E e^{\theta Z} \quad (3.1)$$

$$P(Z \leq z) \leq \inf_{\theta > 0} e^{\theta z} E e^{-\theta Z}. \quad (3.2)$$

Proof of Lemma 1: Given $X(t) = x$, the number N of white balls in boxes $i \geq 1$ which are moved in the next step is

$$N = \sum_{i \geq 1} B_i, \quad (B_i) \text{ independent, } B_i \stackrel{\mathcal{D}}{=} \text{binomial}(x_i, 2^{-i}).$$

In order that some ball be recolored, it is necessary that $N \leq 1$, and so it suffices to prove

$$P(N \leq 1) \leq 2 \exp\left(-\frac{1}{2}f(x)\right).$$

However, this follows from (3.2) with $\theta = \log 2$ via routine calculations.

Proof of Lemma 2: Suppose we modify the process as follows: all balls moved from box L_0 to box $L_0 + 1$ are colored red. For this modified process let $X^*(t) = (X_i^*(t): 1 \leq i \leq L_0)$ be the counts of white balls in boxes i . A routine coupling argument shows we can construct X^* and X together in such a way that the white balls in X^* are identified with a subset of the white balls in X , and so

$$X_i^*(t) \leq X_i(t), \quad i \leq L_0, t \geq 0. \quad (3.3)$$

Now X^* is an irreducible countable-state Markov chain; it is dominated by $Y(t)$ (the process counting balls regardless of color) which has a stationary distribution, and it easily follows that X^* is ergodic. Now define (compare (2.4)) A_s^* to be the event

$$X_i^*(s) \geq \nu 2^i, \quad \text{for all } 1 \leq i \leq L_0.$$

From the coupling, $A_s^* \subset A_s$. Now if we run the process X^* with its stationary distribution, then the event A_s^* has nonzero probability. By ergodicity $P(A_s^* \text{ happens infinitely often}) = 1$ regardless of the initial distribution, and the lemma is proved.

Lemma 6: a) If $Z \stackrel{\mathcal{D}}{=} \text{binomial}(N, q)$, then $P(Z \geq 2Nq) \leq \exp(-(3-e)Nq)$. b) Let (\mathcal{F}_i) be increasing σ -fields, let $F_j \in \mathcal{F}_j$, and let $D_n = \sum_{i=1}^n 1_{F_i}$. Let $\Omega_j = \{P(F_{i+1} | \mathcal{F}_i) \leq q \text{ for all } 0 \leq i < j\}$. Then

$$P(D_N \geq 2Nq, \Omega_N) \leq \exp(-(3-e)Nq).$$

Proof: In case a) we have

$$Ee^{\theta Z} = (1 + (e^\theta - 1)q)^N, \quad \theta > 0,$$

and the result follows from (3.1), setting $\theta = 1$. In case b) we shall prove

$$Ee^{\theta D_n} 1_{\Omega_n} \leq (1 + (e^\theta - 1)q)^n, \quad 1 \leq n \leq N, \quad (3.4)$$

and then the result for b) follows similarly.

To prove (3.4),

$$\begin{aligned} E(e^{\theta D_{n+1}} 1_{\Omega_{n+1}} | \mathcal{F}_n) &= 1_{\Omega_{n+1}} e^{\theta D_n} E(e^{\theta 1_{F_{n+1}}} | \mathcal{F}_n) \\ &\leq 1_{\Omega_{n+1}} e^{\theta D_n} (1 + (e^\theta - 1)q) \end{aligned}$$

because on Ω_{n+1} the conditional probability of F_{n+1} is at most q . Since $\Omega_{n+1} \subset \Omega_n$, taking expectations gives

$$E(1_{\Omega_{n+1}} e^{\theta D_{n+1}}) \leq (1 + (e^\theta - 1)q) E(1_{\Omega_n} e^{\theta D_n}),$$

and (3.4) follows by induction.

Proof of Lemma 3: Fix t and let $u_0 = t - 2^{4+i} \geq t_0$. Then

$$D_i(t) = \sum_{j=u_0}^{t-1} 1_{F_j}$$

where F_j is the event that some ball is recolored on the j th move. Recall $\mathcal{F}_j = \sigma(X(u), Y(u): u \leq j)$. For $u_0 \leq j < t$,

$$\begin{aligned} P(F_{j+1} | \mathcal{F}_j) &\leq 2 \exp\left(-\frac{1}{2}f(X(j))\right) \text{ by Lemma 1} \\ &\leq 2 \exp\left(-\frac{1}{4}\nu L(j)\right) \text{ on } B_{t-1} \\ &\leq 2 \exp\left(-\frac{1}{4}\nu L_0\right) \text{ on } B_{t-1}. \end{aligned}$$

By taking L_0 sufficiently large, we may assume $2 \exp(-(1/4)\nu L_0) \leq 2^{-5}\nu/10 = q$, say. Then applying Lemma 6b),

$$P(D_i(t) \geq 2 \cdot 2^{4+i}q, B_{t-1}) \leq \exp(-(3-e)2^{4+i}q).$$

Substituting the value of q ,

$$P(D_i(t) \geq \nu 2^i/10, B_{t-1}) \leq \exp(-(3-e)2^i/20).$$

Using the hypothesis $i \geq L(t)/10$, this yields the result of Lemma 3 for unconditional probabilities, and the argument works unchanged conditionally on \mathcal{F}_{t_0} .

Proof of Lemma 4: Recall that $Y(t) = (Y_i(t): i \geq 1)$ is the process counting balls regardless of color and is started with its stationary distribution (2.1). Let Y^* be the time-reversed process. Then Y^* is necessarily stationary Markov. It is easy to verify that the transition mechanism for Y^* is the same as for Y , except that balls removed from a box i are put into box $i-1$ instead of box $i+1$. Fix i . Then the quantity $Z_i = Z_i(t)$ of Lemma 4 is distributed as the number of balls in the time-reversed process which start in box i at time 0 and are in box 0 at time 2^{4+i} . Since $Y_i^*(0) \stackrel{\mathcal{D}}{=} \text{Poisson}(\nu 2^i)$,

$$Z_i \stackrel{\mathcal{D}}{=} \text{Poisson}(\nu 2^i \beta) \quad (3.5)$$

where β is the chance that a given ball started in box i in the time-reversed process reaches box 0 before time 2^{4+i} . Now

$$\beta = P(T_i + T_{i-1} + \cdots + T_1 \leq 2^{4+i})$$

where T_j is the holding time in box j . However, $ET_j = 2^j$, so summing over j and using Markov's inequality,

$$\beta \geq 7/8. \quad (3.6)$$

Now routine calculations from (3.2) with $\theta = 1/9$ show that if

$$Z \stackrel{\mathcal{D}}{=} \text{Poisson}(\hat{\lambda}), \quad \hat{\lambda} \geq \lambda,$$

then

$$P\left(Z \leq \frac{8}{9}\lambda\right) \leq \exp\left(-\frac{1}{2}\lambda/9^2\right).$$

Applying (3.5) and (3.6),

$$P\left(Z_i \leq \frac{7}{9}\nu 2^i\right) \leq \exp(-\nu 2^i/200),$$

establishing the lemma.

Proof of Lemma 5: We must estimate a conditional probability given \mathcal{F}_{t_0} on A_{t_0} . Therefore, we may suppose $X(t_0) = x(t_0)$ is given and (by definition of A_{t_0}) that $x_i(t_0) \geq \nu 2^i$, $i \leq L_0$, and it will suffice to prove that

$$P\left(X_i(t) \leq \frac{9}{10}\nu 2^i\right) \leq \exp(-\nu 2^i/200),$$

$$0 \leq t - t_0 \leq 2^{i-5}, i \leq L_0. \quad (3.7)$$

Fix i , and let c be the least integer not less than $\nu 2^i$. Consider a subset of c white balls in box i at time t_0 , and let N be the number of these balls which remain in box i at time $t_0 + 2^{i-5}$. For $t_0 \leq t \leq t_0 + 2^{i-5}$,

$$X_i(t) \geq N \stackrel{\mathcal{D}}{=} \text{binomial}(c, (1 - 2^{-i})^{2^{i-5}}).$$

Thus it suffices to prove

$$P\left(N \leq \frac{9}{10}c\right) \leq \exp(-c/200).$$

Let $\hat{N} \stackrel{\mathcal{D}}{=} \text{binomial}(c, 2^{-5})$. Since

$$(1 - 2^{-i})^{2^{i-5}} \geq 1 - 2^{-5},$$

N stochastically dominates $c - \hat{N}$, and so it suffices to

prove

$$P(\hat{N} \geq c/10) \leq \exp(-c/200).$$

However, this follows from Lemma 6a).

ACKNOWLEDGMENT

My thanks to F. Kelly for suggesting the problem, to W. Rosenkrantz for catching a slip in the first version, and to a referee for helpful comments.

REFERENCES

- [1] L. Breiman, *Probability*. Reading, MA: Addison-Wesley, 1968.
- [2] M. Cottrell, J-C Fort, and G. Malgouyres, "Large deviations and rare events in the study of stochastic algorithms," *IEEE Trans. Automat. Contr.*, vol. AC-28, pp. 907-920, Oct. 1983.
- [3] G. Fayolle, "On ergodicity conditions for special p -dimensional random walks," in *Stochastic Dynamic Optimization and Applications in Scheduling and Related Areas*, F. J. Radermacher, Ed. Passau, W. Germany: Universität Passau, 1985, pp. 23-26.
- [4] J. Goodman, A. G. Greenberg, N. Madras, and P. March, "On the stability of the Ethernet," *Adv. Appl. Probab.*, to appear.
- [5] B. Hajek, "Acknowledgement-based random access transmission control: An equilibrium analysis," in *Proc. Int. Conf. Communications*. New York: IEEE Press, 1982.
- [6] *IEEE Trans. Inform. Theory*, vol. IT-31, Mar. 1985.
- [7] F. P. Kelly, "Stochastic models of computer communication systems," *J. Roy. Statist. Soc. B*, vol. 47, pp. 379-395, 1985.
- [8] F. P. Kelly and I. M. MacPhee, "The number of packets transmitted by collision detect random access schemes," *Ann. Probab.*, to appear.
- [9] W. A. Rosenkrantz, "Some theorems on the instability of the exponential back-off protocol," in *Performance '84*, E. Gelenbe, Ed. Amsterdam, The Netherlands: Elsevier, 1984, pp. 199-205.