1-1-1981

Bounds on the Performance of Protocols for a Multiple-Access Broadcast Channel

Nicholas Pippenger
Harvey Mudd College

Bounds on the Performance of Protocols for a Multiple-Access Broadcast Channel

NICHOLAS PIPPENGER, MEMBER, IEEE

Abstract—**A general model is presented for synchronous protocols that resolve conflicts among message transmissions to a multiple-access broadcast channel. An information-theoretic method is used now to show that if only finitely many types of conflicts can be distinguished by the protocol, utilization of the channel at rates approaching capacity is impossible. A random-coding argument is used to show that if the number of conflicting transmissions can be determined (which requires distinguishing infinitely many types of conflicts) then utilization of the channel at rates arbitrarily close to capacity can be achieved.**

I. Introduction

CONSIDER THE following idealized situation. Messages arrive for transmission at geographically dispersed locations according to a Poisson process in time with rate λ messages per unit time throughout the interval $[0, \mu)$. A multiple-access broadcast channel operates synchronously and is capable of transmitting one message per unit time; consider that the transmissions occur at the "service epochs" $1,2,3,\cdots$. A protocol is used to coordinate the transmission of messages over the channel.

The protocol operates by designating a sequence $Y_1, Y_2, Y_3, \cdots$ of subsets of time. At the service epoch T, an attempt is made to transmit each message that arrived during Y_T and that was not successfully transmitted at some preceding service epoch. There may be no such messages, in which case no transmission occurs. Or there may be just one such message, in which case it is successfully transmitted. Or there may be two or more such messages, in which case they are simultaneously transmitted; this simultaneous transmission causes a "conflict" and none of the messages are successfully transmitted. The subset Y_{T+1} designated for service epoch $T+1$ may depend upon which of these three outcomes occurs at each of the preceding service epochs $1,2,\cdots,T$.

If the expected number of messages arriving during the interval $[0,\mu)$ is denoted $\nu = \lambda\mu$ and the expected number of steps needed to transmit these messages successfully is denoted σ, we may take the ratio ν/σ as a measure of the "throughput" of the protocol. A number of protocols of the type described have been presented in the literature; see, for example, Hayes [8], Capetanakis [2], and Tsybakov and Mikhailov [7]. In her thesis [5], Mosely presents a protocol achieving a throughput of $0.48775\cdots$, which appears to be the highest throughput achieved thus far.

Manuscript received July 17, 1979; revised April 14, 1980.

The author is with the Research Laboratory, IBM Corporation, San Jose, CA 95193.

In Section III. an information-theoretic method will be used to derive a bound on the performance of protocols. It will be shown that

$$\sigma \geq \nu/\xi, \tag{1}$$

for any protocol of the type described, where $\xi = 0.744\cdots$ is the unique solution of the equation

$$-x\log x-(1-x)\log(1-x)+(1-x)\log 2 = x\log e$$

in the range $0 \leq x \leq 1$ ($e = 2.718\cdots$ denotes the base of the natural exponential). This bound supports the following conclusion: for Poisson arrivals, no protocol of the type described can approach throughput one or full utilization of the channel.

It is instructive to determine the basis for this conclusion; that is, how the assumptions concerning the arrival process and the protocol might be changed without altering the conclusion. First it should be noted that the conclusion does not depend particularly on whether time is continuous or discrete. Consider what happens, for example, if the Poisson process is replaced by the outcome of a series of Bernoulli trials. Suppose that messages arrive independently and with the stationary probability p at each of the "arrival epochs" $1,2,3,\cdots$, so that the expected number of messages arriving during $\{1,2,\cdots,M\}$ is $\nu = Mp$. As before, let the expected number of steps needed to transmit these messages successfully be denoted σ. Then a bound of the form (1) can be derived (simply by replacing Poisson probabilities with Bernoulli probabilities in the proof), the constant ξ goes over to a constant ξ_p, which depends on p and is determined by the equation

$$-x\log x-(1-x)\log(1-x)+(1-x)\log 2 = xF(p),$$

where

$$F(p) = -((1-p)\log(1-p))/p.$$

As p increases from zero to one, $F(p)$ decreases from $\log e$ to zero, and ξ_p increases from $\xi = 0.744\cdots$ to one. (See Fig. 1 and 2.) Thus Poisson arrivals are equivalent as usual to numerous Bernoulli arrivals with small probabilities. Furthermore, $\xi_p < 1$ if $p < 1$, and so the conclusion stated for Poisson arrivals also holds for Bernoulli arrivals, provided p is bounded below one. (This shows that the conclusion is not attributable to the fact that the number of Poisson arrivals is potentially unlimited, or to the fact that they may occur arbitrarily closely in time.)

0018-9448/81/0300-0145$00.75

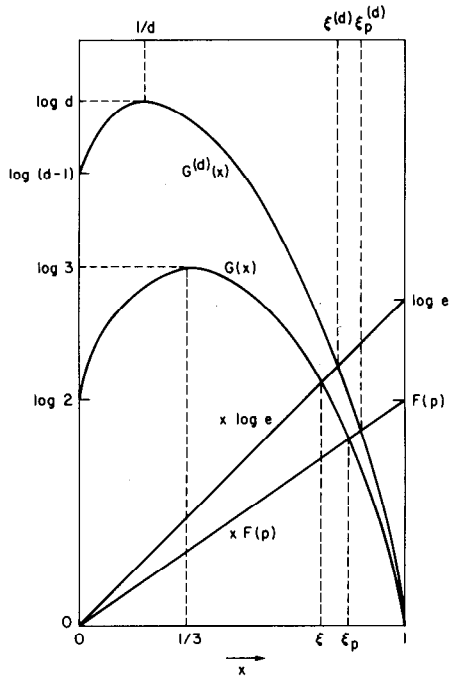

Fig. 1. Numbers ξ, ξ_p, $\xi^{(d)}$, and $\xi_p^{(d)}$ are the abscissae at which the graphs of the functions $G(x) = -x\log x - (1-x)\log(1-x) + (1-x)\log 2$ and $G^{(d)}(x) = -x\log x - (1-x)\log(1-x) + (1-x)\log(d-1)$ intersect those of the linear functions $x\log e$ and $xF(p)$.

Next, it should be noted that the assumption of ternary branching can be relaxed to d-ary branching for any fixed d. This allows consideration of protocols that learn more about the number of simultaneous transmissions at each step than is expressed by the three cases $0, 1, 2$ or more. They might learn for example, what is expressed by the d cases $0, 1, \cdots, d-2, d-1$ or more. Bounds of the form (1) can again be derived; for Poisson arrivals the constant $\xi^{(d)}$ is determined by the equation

$$-x\log x - (1-x)\log(1-x) + (1-x)\log(d-1) = x\log e,$$

while for Bernoulli arrivals the constant $\xi_p^{(d)}$ is determined by the equation

$$-x\log x - (1-x)\log(1-x) + (1-x)\log(d-1) = xF(p),$$

where $F(p)$ is as defined above. (See Figs. 1 and 2.) For finite d, $\xi^{(d)} < 1$, and $\xi_p^{(d)} < 1$ if $p < 1$. Thus the conclusion stated for ternary branching holds also for d-ary branching, provided d is bounded.

The requirement that p be bounded below one for discrete time is obviously necessary, for as p approaches one, the arrivals become completely predictable, and the protocol that designates the singletons seriatim approaches full utilization of the channel. The requirement that d be bounded is also necessary; as d tends to infinity, protocols with d-ary branching can approach full utilization of the channel. The limiting case of this phenomenon is presented in Section IV, where a random-coding argument is used to show that a protocol with "infinitary branching" (that is, one that learns the number of simultaneous transmissions

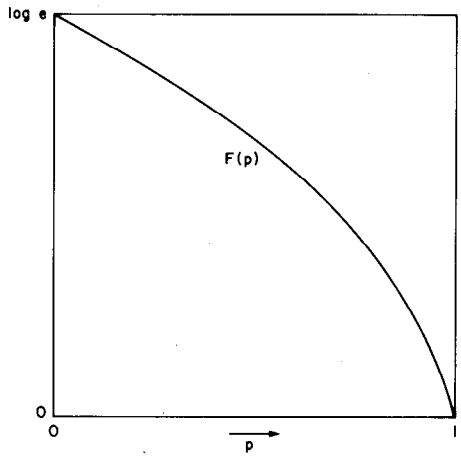

Fig. 2. Function $F(p) = -((1-p)\log(1-p))/p$ decreases from $\log e$ to zero as p increases from zero to one; it is concave and satisfies $F'(0) = 0$, $F'(1) = -\infty$.

at each step) can achieve

$$\sigma \leqslant \nu + O\big(\nu/(\log \nu)^{1/2}\big). \tag{2}$$

This will be shown for Poisson arrivals; the same result can be derived for Bernoulli arrivals.

The protocol that achieves this bound can be adapted to give, for any $\eta < 1$, a protocol with d-ary branching (for any sufficiently large d) for which $\sigma \leqslant \nu/\eta$ (for all sufficiently large ν). No claims for its practicality are intended, however: the protocol does not correspond to an easily implemented algorithm, and the constant implicit in the $O(\nu/(\log \nu)^{1/2})$ term is large.

Finally, it should be mentioned that there is another class of protocols that are superficially dissimilar from those described above, but nonetheless amenable to the same analysis. These are protocols that use independent randomization at the geographically dispersed locations instead of, or in addition to, the arrival times of the messages to determine which message to transmit at each service epoch. (See Abramson [1] for examples and references.) As will become clear in the proofs, however, what matters is that messages can be separated from one another by the outcome of some random process; it is immaterial whether this is the result of differences in random arrival times, explicit randomization, or some combination of the two.

II. Model

Let X denote a finite interval of time and let the random variable $\boldsymbol{E}$ denote a set of Poisson arrivals during X. With probability one, $\boldsymbol{E}$ will be a finite subset of X; its elements will be called *messages*. In the Introduction, X was taken to be $[0, \mu)$ and the arrival rate was taken to be λ; here it will be equivalent and simpler to take X to be $[0, 1)$ and the arrival rate to be $\nu = \mu\lambda$.

A *protocol* for X will be modeled by an infinite tree in which there is an initial node called the root, in which each node K is connected by branches to three offspring $K^{(0)}$, $K^{(1)}$, $K^{(2)}$ (which may be either leaves or other nodes) and

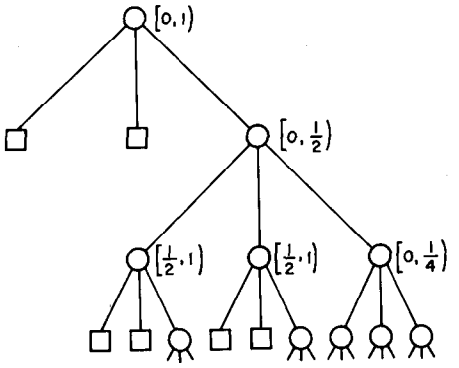

Fig. 3. Part of a protocol with ternary branching for [0,1). Nodes are indicated by circles, with the root at the top; leaves are indicated by squares.

in which each node K is labeled with a measurable subset $Y(K)$ of X. (See Fig. 3.)

The *execution* of a protocol for X with respect to a finite subset E of X is a path through the tree defined as follows. Let K_0, the first node of the path, be the root, and let $E_0 = E$. Suppose that K_R and E_R have been determined; then K_{R+1} is defined to be $K_R^{(0)}$, $K_R^{(1)}$, or $K_R^{(2)}$ according as $E_R \cap Y(K_R)$ contains 0, 1, or 2 or more messages. If $E_R \cap Y(K_R)$ contains just one message, this message is said to be *transmitted successfully* by K_R, and E_{R+1} is obtained from E_R by deleting this message; otherwise $E_{R+1} = E_R$.

A protocol for X will be called *valid* if, for every finite subset E of X, the execution $K_0, K_1, K_2, \cdots$ of the protocol with respect to E terminates after finitely many steps at a leaf $K_S = L$ after each message in E has been successfully transmitted (so that E_S is empty). The execution of a valid protocol with respect to the random set $\boldsymbol{E}$ is a random path $\boldsymbol{K}_0, \boldsymbol{K}_1, \boldsymbol{K}_2, \cdots$ which terminates at a random leaf $\boldsymbol{L}$ after a random number $\boldsymbol{S}$ of steps during which a random number $\boldsymbol{N}$ of messages are successfully transmitted.

Let $p(S)$ denote the probability that $\boldsymbol{S}=S$; then

$$\sum_S p(S)=1,$$

and

$$\sigma=\sum_S p(S)S$$

is the expected number of steps in the execution of the protocol with respect to $\boldsymbol{E}$. Let $p(N)$ denote the probability that $\boldsymbol{N}=N$. Then

$$\sum_N p(N)=1, \tag{3}$$

and

$$\nu=\sum_N p(N)N \tag{4}$$

is the expected number of messages in $\boldsymbol{E}$.

III. Proof of (1)

Before launching into the proof, it may be helpful to give the gist of it. In addition to generating messages, the arrival process also generates entropy in the form of uncertainty as to the locations of the messages in time. A protocol must resolve some of this uncertainty in order for the messages to be transmitted successfully. There are two ways in which no message can be successfully transmitted (no transmissions and two or more transmissions), so the protocol can gain up to one bit of information in this case. There is, however, only one way in which a message can be successfully transmitted, so the protocol gains no additional information in this case beyond the fact that a message was successfully transmitted. In this sense, a protocol can learn more from failure than from success, and it must risk failures to gain the required information. With finitary branching, only a bounded amount of information can be gained in return for such a risk. The final bound reflects the compromise necessary between the desire to transmit a message successfully at a given step and the desire to gain information in order to transmit messages successfully at later steps.

Consider a function that assigns to every subset E of X containing N messages a partition U of X into N measurable blocks $U_1, \cdots, U_N$ such that each block in U contains just one message in E. A random partition $\boldsymbol{U}$ obtained in this way from the random set $\boldsymbol{E}$ will be called a *resolution* of $\boldsymbol{E}$.

Let $\boldsymbol{U}$ be a resolution of $\boldsymbol{E}$ and let $p(U)$ denote the probability that $\boldsymbol{U}=U$. Then

$$\sum_U p(U)=1.$$

Let

$$\eta=-\sum_U p(U)\log p(U) \tag{5}$$

denote the entropy of this distribution. Then

$$\eta \geqslant \nu \log e. \tag{6}$$

To see this, suppose that U is a partition of X into N blocks $U_1, \cdots, U_N$, and that $u_1, \cdots, u_N$ are the measures of these blocks, so that

$$\sum_{1 \leqslant M \leqslant N} u_M = 1.$$

If $\boldsymbol{U}=U$, each block of U must contain just one message from $\boldsymbol{E}$; thus

$$p(U) \leqslant \prod_{1 \leqslant M \leqslant N} (u_M \nu \exp -(u_M \nu))$$

$$= \Big(\prod_{1 \leqslant M \leqslant N} u_M \Big) \nu^N \exp -\nu,$$

where "exp" denotes the natural exponential. Since a geometric mean is bounded above by the corresponding arithmetic mean,

$$p(U) \leqslant \Big(\sum_{1 \leqslant M \leqslant N} u_M / N \Big)^N \nu^N \exp -\nu$$

$$= (\nu/N)^N \exp -\nu.$$

Substituting this bound for the argument of the logarithm

in (5) and rearranging yields

$$\eta \geq \nu \log e + \sum_N p(N) N \log(N/\nu).$$

Since $x\log(x/\nu)$ is a convex function of x, (3) and (4) imply that

$$\sum_N p(N) N \log(N/\nu) \geq \nu \log(\nu/\nu) = 0.$$

This completes the proof of (6).

Consider now a particular valid (but otherwise arbitrary) protocol. For each leaf L, let $p(L)$ denote the probability that $\boldsymbol{L} = L$, that is, that execution terminates at L. Then

$$\sum_L p(L) = 1.$$

Let

$$\zeta = -\sum_L p(L) \log p(L) \tag{7}$$

denote the entropy of this distribution. Then

$$\zeta \geq \nu \log e. \tag{8}$$

To see this, observe that a valid protocol determines a resolution of $\boldsymbol{E}$ in the following way. Suppose that the execution with respect to E is a path $K_0, K_1, K_2, \cdots$ which terminates at the leaf L, and suppose that $K_{R(1)}, \cdots, K_{R(N)}$ are the nodes at which the messages in E are successfully transmitted. Let $U_1 = Y(K_{R(1)})$, $U_2 = Y(K_{R(2)}) - U_1, \cdots, U_{N-1} = Y(K_{R(N-1)}) - (U_1 \cup \cdots \cup U_{N-2})$, $U_N = X - (U_1 \cup \cdots \cup U_{N-1})$, and let $U = \{U_1, \cdots, U_N\}$. It is easy to see that this defines a resolution $\boldsymbol{U}$ of $\boldsymbol{E}$. Since L determines U, the entropy ζ of $\boldsymbol{L}$ is at least as large as the entropy η of $\boldsymbol{U}$. This completes the proof of (8).

For each node K, let $p(K)$ denote the probability that K is a member of $\{\boldsymbol{K}_0, \boldsymbol{K}_1, \boldsymbol{K}_2, \cdots\}$, that is, that execution passes through K. Then

$$\sum_K p(K) = \sigma, \tag{9}$$

since execution passes through a node of the protocol at each step.

For each node K, let $q(K,0)$, $q(K,1)$, and $q(K,2)$ denote the conditional probabilities that execution passes from K to $K^{(0)}$, $K^{(1)}$, and $K^{(2)}$, given that execution passes through K. Then $p(K)q(K,1)$ is the probability that execution passes through both K and $K^{(1)}$, and

$$\sum_K p(K) q(K,1) = \nu, \tag{10}$$

since execution passes from a node K to its offspring $K^{(1)}$ when and only when a message is successfully transmitted by K.

For each node K,

$$\sum_{0 \leq J \leq 2} q(K,J) = 1.$$

Let

$$h(K) = -\sum_{0 \leq J \leq 2} q(K,J) \log q(K,J)$$

denote the entropy of this distribution. Each probability $p(K)$ or $p(L)$ can be expressed as a product of conditional probabilities $q(K_0, J_0) \cdots q(K_R, J_R)$, where $K_0, \cdots, K_R$ is the path from the root to K or L, and $J_0, \cdots, J_R$ designate the branches taken along this path. This allows (7) to be rewritten as

$$\zeta = \sum_K p(K) h(K).$$

Thus

$$\sum_K p(K) h(K) \geq \nu \log e.$$

Let

$$G(x) = -x\log x - (1-x)\log(1-x) + (1-x)\log 2$$

denote the entropy of the scheme $((1-x)/2, x, (1-x)/2)$. Then

$$h(K) \leq G(q(K,1)),$$

since replacing $q(K,0)$ and $q(K,2)$ by their arithmetic mean $(1 - q(K,1))/2$ can only increase the entropy. Thus

$$\sum_K p(K) G(q(K,1)) \geq \nu \log e. \tag{11}$$

The conclusion is now at hand. Since $G(x)$ is a concave function of x, (9), (10), and (11) imply that

$$\sigma G(\nu/\sigma) \geq \nu \log e.$$

The inequality $G(x) \geq x \log e$ holds if and only if $x \leq \xi$, where $\xi = 0.744\cdots$ is the unique solution of $G(x) = x \log e$ in the range $0 \leq x \leq 1$. This completes the proof of (1).

IV. Proof of (2)

Again, it may be helpful to give the gist of the proof. With infinitary branching, the possibility exists of gaining an unbounded amount of information at a single step. To exploit this possibility, the protocol must designate large subsets, which are likely to contain many messages, so that the entropy of the number of messages will be large. These subsets must overlap in complicated ways, so that new information is gained at each step; complicated patterns of overlap can be obtained in a manageable way by means of a random-coding argument. While information is being gathered, successful transmissions of messages only confuse matters since they change the numbers of messages in various subsets. The protocol can avoid successful transmissions by including large subsets that are known to contain at least two messages as "ballast" in the subsets it designates, thus deliberately causing a conflict. Once enough information has been gathered, all the messages can be successfully transmitted with few wasted steps. The final bound again reflects a compromise between gaining information and successfully transmitting messages.

The notation $O(f(x))$ will be used (following Bachmann) to denote a function of x (not necessarily the same function at every occurrence) bounded above in absolute value by $cf(x)$ for some constant $c>0$. The notation $\Omega(f(x))$ will be defined similarly (following Knuth), with "bounded above in absolute value" replaced by "bounded below."

The protocol presented in this section will conform to the model presented in Section II, with ternary branching replaced by infinitary branching. This means simply that each node K is connected by branches to infinitely many offspring $K^{(0)}, K^{(1)}, K^{(2)}, \cdots$, and that K_{R+1} is defined to be $K_R^{(0)}, K_R^{(1)}, K_R^{(2)}, \cdots$ according as $E_R \cap Y(K_R)$ contains $0, 1, 2, \cdots$ messages. The protocol will be presented informally, but it should be clear that it could be formalized within this model.

With a single step, the protocol determines the number $\boldsymbol{N}$ of messages in $\boldsymbol{E}$. In this case then, if $\boldsymbol{N}=N\leqslant 1$, any message in $\boldsymbol{E}$ is thereby successfully transmitted. If $\boldsymbol{N}=N \geqslant 2$, a simple calculation shows that the conditional distribution of the N messages is uniform and independent throughout X. It will be sufficient to show that the task of the protocol can be completed with an expected number $N+O(N/(\log N)^{1/2})$ of steps, for (2) will then follow by averaging:

$$\sum_N \left\{N+O\left(N/(\log N)^{1/2}\right)\right\}(\nu^N/N!)\exp-\nu \leqslant \nu+O\left(\nu/(\log \nu)^{1/2}\right).$$

Suppose then that $N\geqslant 2$ messages are distributed uniformly and independently throughout $X=[0,1)$. Let

$$B=2\left\lceil N(\log N)^{1/2}\right\rceil,$$
$$\epsilon=1/B,$$

and divide X into B subintervals, each of length ϵ that will be called *blocks*: $X_1=[0,\epsilon),\cdots,X_B=[1-\epsilon,1)$. Let $\boldsymbol{N}_1,\cdots,\boldsymbol{N}_B$ denote the random numbers of messages in $\boldsymbol{E}$ falling in the blocks $X_1,\cdots,X_B$.

The remaining action of the protocol will be divided into two phases. The task of the first phase will be to determine $\boldsymbol{N}_1,\cdots,\boldsymbol{N}_B$. This could obviously be accomplished with B steps. The crux of the proof will be to show that it can be accomplished with an expected number $O(N/(\log N)^{1/2})$ of steps. The first phase will satisfy the following condition: if any message in $\boldsymbol{E}$ is successfully transmitted, then all messages in $\boldsymbol{E}$ are successfully transmitted. Thus the first phase either acts purely as a gatherer of information, or it completes the task of the protocol and eliminates the need for a second phase. The second phase, if necessary, will complete the task of the protocol; this will be done with an expected number $N+O(N/(\log N)^{1/2})$ of steps.

An essential ingredient of both phases will be a crude procedure for causing a known number $M\geqslant 2$ of messages that are uniformly and independently distributed throughout an interval to be transmitted successfully with an expected number $O(M)$ of steps. This procedure, which will be called *recursive binary splitting*, is as follows. With two steps, the protocol determines the numbers M_1 and M_2 of messages in the left and right halves of the interval. If $M_1\leqslant 1$ or $M_2\leqslant 1$ or both, any messages in the corresponding subintervals are thereby successfully transmitted. If $M_1\geqslant 2$ or $M_2\geqslant 2$ or both, the protocol causes the messages in the corresponding subintervals to be transmitted successfully by recursive binary splitting.

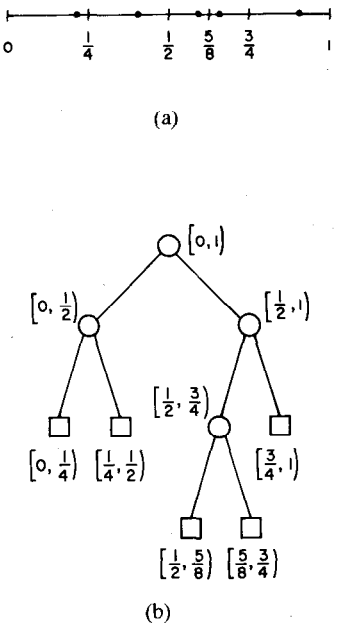

Fig. 4. (a) Five messages in the interval [0,1). (b) Binary tree corresponding to application of recursive binary splitting to these messages. This type of tree, which corresponds to a single execution of a protocol, should not be confused with the type shown in Fig. 3, which represents a protocol in its entirety.

This procedure is egregiously inefficient but has the merit of being easy to describe and analyze. Application of the procedure to the messages in an interval (which may be taken to be [0,1) without loss of generality) gives rise to a binary tree in which the root corresponds to the interval [0,1), the nodes correspond to dyadic subintervals (that is, subintervals of the form $[(K-1)/2^d, K/2^d)$, for $1\leqslant K\leqslant 2^d$) which contain two or more messages, and the leaves correspond to dyadic subintervals which contain at most one message. (See Fig. 4.) The number of steps performed by the procedure is one less than the number of nodes and leaves in the tree since there is a step for each node or leaf except the root. Thus the number of steps is twice the number of nodes since in any binary tree the number of nodes is one less than the number of leaves. There are 2^d dyadic subintervals of length $1/2^d$, and the probability that such a subinterval contains two or more messages is

$$\sum_{2\leqslant J\leqslant M}\binom{M}{J}(1/2^d)^J(1-1/2^d)^{M-J}.$$

Thus the expected number of steps is

$$2\sum_d 2^d \sum_{2\leqslant J\leqslant M}\binom{M}{J}(1/2^d)^J(1-1/2^d)^{M-J}.$$

This expression is $O(M)$, as is easily seen by considering separately the terms for which $2^d \leq 2M$ and those for which $2^d > 2M$: for the former, the inner sum is at most one, and the contribution of these terms to the expression is at most $8M$. For the latter, the inner sum is at most $(M/2^d)^2$, and the contribution of these terms is at most $2M$. (A more profound analysis of this expression follows the realization that recursive binary splitting of M messages in the interval $[0,1)$ is equivalent to radix-exchange sorting of M records with keys in this interval, which has been analyzed by Knuth [4, section 5.2.2]. The expression is $2A_M$ in the notation of that analysis, where it is shown to be approximately $(2\log_2 e)M$. Knuth's analysis implicitly reduces the uniform and independent distribution of the keys to a Poisson distribution, thus reversing the strategy of the present proof.)

The key to the operation of the first phase will be the following combinatorial proposition. For every natural number L, there is a natural number $K = O(L/\log L)$ and a $K \times L$ binary matrix F with the following property: each L-dimensional vector G of natural numbers $G_1, \cdots, G_L$ satisfying $G_1 + \cdots + G_L \leq L$ is uniquely determined by the corresponding K-dimensional vector FG. This proposition will be proved by a random-coding argument. (The proposition and its proof are generalizations of Theorem 1 and its proof in the paper [3] by Erdös and Rényi.)

For brevity, an L-dimensional vector G of natural numbers $G_1, \cdots, G_L$ satisfying $G_1 + \cdots + G_L \leq L$ will be called an *L-composition*. With this terminology, the desired property of F is that if G and G' are distinct L-compositions, then FG and FG' are distinct.

Two L-compositions G and G' will be called *disjoint* if the sets $\mathcal{G} = \{J\colon G_J > 0\}$ and $\mathcal{G}' = \{J\colon G'_J > 0\}$ of indices for which their components do not vanish are disjoint. If $G_1, \cdots, G_L$ and $G'_1, \cdots, G'_L$ are distinct L-compositions for which $FG = FG'$, and if $H_1 = \min\{G_1, G'_1\}, \cdots, H_L = \min\{G_L, G'_L\}$, then $G - H$ and $G' - H$ are distinct and disjoint L-compositions for which $F(G-H) = F(G'-H)$. Thus an equivalent formulation of the desired property is that if G and G' are distinct and disjoint L-compositions, then FG and FG' are distinct.

A pair of distinct and disjoint L-compositions G and G' for which $\mathcal{G}$ contains Q indices and $\mathcal{G}'$ contains Q' indices will be called a *failure mode* of size $Q + Q'$. The size R of a failure mode must be in the range $1 \leq R \leq L$. The number of failure modes of size R is

$$\sum_{Q+Q'=R} \binom{L}{QQ'}\binom{L}{Q}\binom{L}{Q'},$$

since the multinomial coefficient $\binom{L}{QQ'}$ enumerates the ways of choosing the sets $\mathcal{G}$ and $\mathcal{G}'$, $\binom{L}{Q}$ enumerates the ways of choosing G in accordance with $\mathcal{G}$, and $\binom{L}{Q'}$ enumerates the ways of choosing G' in accordance with $\mathcal{G}'$. A rough bound for this expression in terms of L and R is

$$\begin{aligned}\sum_{Q+Q'=R} \binom{L}{QQ'}\binom{L}{Q}\binom{L}{Q'} &\leq \binom{L}{R}2^R \sum_{Q+Q'=R} \binom{L}{Q}\binom{L}{Q'} \\ &= \binom{L}{R}2^R \binom{2L}{R} = \exp O(R + R\log(L/R)).\end{aligned}$$

Let $\boldsymbol{F}$ be a random $K \times L$ binary matrix, in which each entry is zero or one equiprobably and independently. It will be shown that for a certain choice of $K = O(L/\log L)$, the probability that $\boldsymbol{F}$ assumes a value F that fails to have the desired property tends to zero as L tends to infinity. It will follow that for a certain choice of $K = O(L/\log L)$, there exists a matrix F with the desired property. (The probabilities in this argument are based on the randomness of $\boldsymbol{F}$; they have nothing to do with probabilities elsewhere in this paper, which are based on the randomness of $\boldsymbol{E}$.)

If F is a $K \times L$ binary matrix, let $\mathcal{F}_1 = \{J\colon F_{1,J} = 1\}, \cdots, \mathcal{F}_K = \{J\colon F_{K,J} = 1\}$ be the sets of column indices for which the entries in the various rows do not vanish. Let G and G' be a failure mode of size R. The number of $K \times L$ binary matrices F for which $FG = FG'$ is at most

$$\left[2^{L-R}\binom{R}{\lfloor R/2 \rfloor}\right]^K,$$

since for each of the K row indices I, there are 2^{L-R} ways of choosing the $L-R$ entries $F_{I,J}$ with column indices J not in $\mathcal{G} \cup \mathcal{G}'$ and at most $\binom{R}{\lfloor R/2 \rfloor}$ ways of choosing the R entries with column indices in $\mathcal{G} \cup \mathcal{G}'$. (Each solution of the equation

$$\sum_{J \in \mathcal{G}} F_{I,J}G_J = \sum_{J \in \mathcal{G}'} F_{I,J}G'_J$$

corresponds to a subset $(\mathcal{G} \cap \mathcal{F}_I) \cup (\mathcal{G}' - \mathcal{F}_I)$ of $\mathcal{G} \cup \mathcal{G}'$. Any pair of distinct solutions corresponds to subsets that are incomparable in the sense that neither is contained within the other. By a theorem of Sperner [6], a family of pairwise incomparable subsets of an R element set can contain at most $\binom{R}{\lfloor R/2 \rfloor}$ subsets.) Thus for a given failure mode G and G' of size R, the probability that $\boldsymbol{FG} = \boldsymbol{FG}'$ is at most

$$\left[\binom{R}{\lfloor R/2 \rfloor}/2^R\right]^K = \exp -\Omega(K\log(R+1)).$$

Combining the estimates derived for the number of failure modes and for the probability of failure for each mode, the total probability of failure (the probability that $\boldsymbol{F}$ does not have the desired property) is at most

$$\sum_{1 \leq R \leq L} \exp\{O(R + R\log(L/R)) - \Omega(K\log(R+1))\}.$$

If

$$K = \lceil CL/\log L \rceil,$$

where C is a sufficiently large constant, then the probability of failure will tend to zero as L tends to infinity. To see

this, consider separately the terms with $1 \leqslant R \leqslant L/(\log L)^2$ and $L/(\log L)^2 < R \leqslant L$. For the former, $O(R + R\log(L/R)) = O(L\log\log L/(\log L)^2)$ and $\Omega(K\log(R+1)) = \Omega(CL/\log L)$; for the latter, $O(R+R\log(L/R)) = O(L)$ and $\Omega(K\log(R+1)) = \Omega(CL)$. This completes the proof of the combinatorial proposition.

In the first phase, the protocol determines with two steps the populations $N^{(1)} = N_1 + \cdots + N_{B/2}$ and $N^{(2)} = N_{B/2+1} + \cdots + N_B$ of $X^{(1)} = X_1 \cup \cdots \cup X_{B/2}$ and $X^{(2)} = X_{B/2+1} \cup \cdots \cup X_B$ (recall that B was chosen to be even). If $N^{(1)} \leqslant 1$ and $N^{(2)} \leqslant 1$, the task of the protocol is complete.

If $N^{(1)} \geqslant 2$ and $N^{(2)} \leqslant 1$ (or $N^{(1)} \leqslant 1$ and $N^{(2)} \geqslant 2$), the protocol applies recursive binary splitting to $X^{(1)}$ (or $X^{(2)}$). In these cases, which arise with probability $O(N2^{-N})$, the task of the protocol is completed with an expected number $O(N)$ of steps.

Finally, if $N^{(1)} \geqslant 2$ and $N^{(2)} \geqslant 2$, the protocol will determine $N_1, \cdots, N_B$ without causing any messages to be transmitted successfully. To do this, it will use the combinatorial proposition with $L = B/2 = O(N(\log N)^{1/2})$, so that $K = O(N/(\log N)^{1/2})$. The $K \times L$ binary matrix F corresponds to K subsets $\mathcal{F}_1, \cdots, \mathcal{F}_K$ of $\{1, \cdots, L\}$. With K steps designating the subsets

$$\left(\bigcup_{J \in \mathcal{F}_1} X_J\right) \cup X^{(2)}, \cdots, \left(\bigcup_{J \in \mathcal{F}_K} X_J\right) \cup X^{(2)},$$

the protocol determines the sums

$$\left(\sum_{1 \leqslant J \leqslant L} F_{1,J} N_J\right) + N^{(2)}, \cdots, \left(\sum_{1 \leqslant J \leqslant L} F_{K,J} N_J\right) + N^{(2)}.$$

Since $N^{(2)} \geqslant 2$, this cannot cause the successful transmission of any messages, and since $N^{(2)}$ is known, it determines the sums

$$\sum_{1 \leqslant J \leqslant L} F_{1,J} N_J, \cdots, \sum_{1 \leqslant J \leqslant L} F_{K,J} N_J.$$

By the combinatorial proposition, this determines $N_1, \cdots, N_L$, since $N_1 + \cdots + N_L \leqslant N \leqslant L$. With K more steps designating the subsets

$$X^{(1)} \cup \left(\bigcup_{J \in \mathcal{F}_1} X_{L+J}\right), \cdots, X^{(1)} \cup \left(\bigcup_{J \in \mathcal{F}_K} X_{L+J}\right),$$

the protocol determines $N_{L+1}, \cdots, N_{2L}$. In this case, then the protocol determines $N_1, \cdots, N_B$ with $2K = O(N/(\log N)^{1/2})$ steps. The expected number of steps performed by the first phase is

$$2 + O(N2^{-N})O(N) + O\left(N/(\log N)^{1/2}\right) = O\left(N/(\log N)^{1/2}\right).$$

In the second phase, if it is necessary, the protocol designates each block X_A for which $N_A = 1$ $(1 \leqslant A \leqslant B)$, thereby causing the messages in such blocks to be successfully transmitted. It then applies recursive binary splitting to each block X_A for which $N_A \geqslant 2$ $(1 \leqslant A \leqslant B)$, thus completing its task. In estimating the expected number of steps performed by the second phase, it may be assumed without undue optimism that the first phase always determines $N_1, \cdots, N_B$ without causing any messages to be transmitted successfully, since the fact that the first phase sometimes completes the task of the protocol, can only decrease the quantity being estimated. There are B blocks. For each block X_A $(1 \leqslant A \leqslant B)$, the probability that $N_A = 1$ is

$$N\epsilon(1-\epsilon)^{N-1} \leqslant N\epsilon,$$

and the protocol performs one step for each such block; the probability that $N_A = M \geqslant 2$ is

$$\binom{N}{M}\epsilon^M(1-\epsilon)^{N-M} \leqslant (N\epsilon)^M,$$

and the protocol performs an expected number $O(M)$ of steps for each such block. Thus the expected number of steps performed by the second phase is at most

$$\begin{aligned} B\left\{N\epsilon + \sum_{2 \leqslant M \leqslant N} (N\epsilon)^M O(M)\right\} &= B\left\{N\epsilon + O(N^2\epsilon^2)\right\} \\ &= N + O\left(N/(\log N)^{1/2}\right). \end{aligned}$$

This completes the proof of (2).

References

[1] N. Abramson, "The throughput of packet broadcasting channels," *IEEE Trans. Commun.*, vol. COM-25, pp. 117–128, Jan. 1977.

[2] J. I. Capetanakis, "Tree algorithms for packet broadcast channels," *IEEE Trans. Inform. Theory*, vol. IT-25, pp. 505–515, Nov. 1979.

[3] P. Erdös and A. Rényi, "On two problems of information theory," *Publ. Hung. Acad. Sci.*, vol. 8, pp. 241–254, 1963.

[4] D. E. Knuth, *The Art of Computer Programming*, vol. 3, *Sorting and Searching*. Reading: MA: Addison-Wesley, 1973.

[5] J. Mosely, "An efficient contention resolution algorithm for multiple access channels," M.S. thesis, Dept. Elec. Eng. and Comp. Sci., Massachusetts Inst. Tech., Cambridge, May 1979.

[6] E. Sperner, "Ein Satz über Untermengen einer endlichen Menge," *Math. Z.*, vol. 27, pp. 544–548, 1928.

[7] B. S. Tsybakov and V. A. Mikhailov, "Free synchronous packet access in a broadcast channel with feedback," *Probl. Inform. Trans.*, vol. 14, pp. 259–280, 1979.

[8] J. F. Hayes, "An adaptive technique for local distribution," *IEEE Trans. Commun.*, vol. COM-26, pp. 1178–1186, Aug. 1978.