

On the Capacity of Infinite Population Multiple Access Protocols

MART L. MOLLE, STUDENT MEMBER, IEEE

Abstract—We present bounds on the maximum channel utilization (with finite average delay) of synchronous multiple access communications protocols serving an infinite population of homogeneous stations. Messages arrive to the system as a series of independent Bernoulli trials in discrete time, with probability p of an arrival at each arrival point (the Poisson limit is explicitly included) and are then randomly distributed among the stations. Pippenger showed that the channel utilization cannot exceed ξ_p , where $\xi_1 = 1$ and $\lim_{p \rightarrow 0} \xi_p \approx 0.744$. Using a “helpful genie” argument, we find the exact capacity for all $p \geq 0.568$ (where we find optimal protocols that obey first-come first-served); for smaller values of p , we present an improved upper bound that decreases monotonically to ≈ 0.6731 in the Poisson limit as $p \rightarrow 0$.

I. INTRODUCTION

CONSIDER a distributed population of stations using a “synchronous” multiple access protocol to exchange messages over a noiseless communications channel. That is, all message transmissions are synchronized to fall into constant length *slots* of duration equal to a message transmission time. The transmissions in a slot can have three outcomes: an *empty slot*, when no station transmits; a *success*, when exactly one station transmits; and a *collision*, when two or more stations transmit simultaneously. We assume that at the end of each slot all stations receive an acknowledgment whenever a message is sent successfully and a nonacknowledgment whenever there is a collision, but we require that the protocol operate without any additional exchange of information. However, the protocol is free to use the history of activity on the channel.

Without loss of generality, such a multiple access protocol may be described as a discrete time Markov process, since an arbitrary amount of information may be encoded into the state description (including any channel history information that affects the behavior of the protocol, and the current backlog of unserved arrivals to the system). At the beginning of each slot, the protocol enters some state and immediately *enables* (i.e., grants transmission rights to) a subset of the stations, the choice of which subset to enable depending on the current state. Each enabled station transmits a message if it has one. All other

stations remain silent for the duration of the slot. At the end of each slot, the next state is chosen as a function of the current state and the activity (i.e., idle, success, or collision) in the current slot. Thus service becomes a probabilistic event. Each attempt to offer service may either result in a successful transmission if exactly one busy station was enabled, or merely provide some information about the distribution of messages among the stations. The goal in defining a protocol is to choose the sequence of enabled subsets so that few slots are wasted as idle slots or collisions.

We say that a protocol is *stable* if a stationary probability distribution exists for the Markov process, i.e., if the backlog of unserved arrivals to the system remains finite with probability one, and we define *capacity* to be the supremum over all arrival rates such that the protocol is stable. For example, with a *finite* number of buffered stations, Tsybakov and Mikhailov [2] have shown that $N(t) \triangleq (n_1(t), n_2(t), \dots, n_M(t))$, where $n_i(t)$ is the queue length at the i th station at time t , is a suitable state description for the (memoryless) ALOHA protocol [3]. Among other results, they proved that the capacity of ALOHA is at least e^{-1} by showing the existence of parameters for which $N(t)$ is ergodic whenever the sum of the station arrival rates is less than e^{-1} .

The complexity of the state space makes the performance analysis of multiple access protocols a difficult task in general. However, determining the capacity of an optimal protocol (in the sense that it attains maximum channel utilization) without explicitly calculating message delays is a much simpler task. A protocol with a larger backlog can always simulate the behavior of a protocol with a smaller backlog by restricting its choice of enabled sets. Thus, to upper-bound the capacity of an optimal protocol, we may assume that the protocol is free to choose enabled sets of any size at any time. Here we may study the behavior of an optimal protocol even when each state encodes only the channel history information that affects the behavior of the protocol (and not the exact value of the backlog), thereby greatly simplifying the state space. We call this simplified process a *collision resolution algorithm*. Therefore, to bound the capacity of an optimal protocol, we need only bound the maximum attainable channel utilization by a collision resolution algorithm.

At present there is considerable theoretical interest in multiple access protocols suitable for an *infinite* number of stations. Infinite population protocols can be used without

Manuscript received March 10, 1980; revised April 27, 1981. This research was supported in part by the Advanced Research Projects Agency of the Department of Defense under Contract MDA 903-77-C-0272, and in part by the Regents of the University of California through a Chancellor's Intern Fellowship.

The author was with the Computer Science Department, University of California, Los Angeles, CA. He is now with the Computer Systems Research Group, University of Toronto, Toronto, ON, Canada M5S 1A1.

modification when there is a finite number of stations. Thus, infinite population protocols are also of considerable practical interest for systems with a large (but finite) number of stations. Their performance in the finite population case is at least as good as in the infinite population case. They are robust in the sense that each station need not be aware of the addresses of every other station in the system, or even of the exact number of stations in the system. This simplifies the procedure by which stations can enter or leave a working system. Furthermore, when infinite population protocols are used, the average message delay for a given channel utilization is insensitive to even large changes in the number of stations sharing the channel.

Unfortunately many seemingly reasonable protocols are infeasible with an infinite number of stations. Round-robin time division multiple access (TDMA), for example, has zero capacity in the infinite population case, for if each of an infinite number of stations had a positive average queue length then the average number in system (and hence the average time in system by Little's result [4]) would be infinite. Consequently, an infinite population multiple access protocol cannot depend on individual *stations* having a queue of messages. Instead, such protocols must search for *messages* directly by, say, examining the time axis for arrivals to the system.

Let us assume that messages may arrive at a series of independent *arrival points*, each having probability p of containing exactly one message and $1 - p$ of containing no messages. (When a message arrives to the system, we assume that it is assigned to one station at random. A protocol as described above is subsequently used by this station to determine, at each slot, whether or not the arrival point for this message has been enabled and hence that the message should be transmitted.) Initially, we say that each arrival point is *unexamined*. As the protocol operates, we may deduce information about the state of an arrival point from the history of channel activity. We will say that an arrival point is *busy* if it becomes known to the protocol that it contains a message that has not yet been successfully transmitted, and *idle* if it becomes known that it does not contain a message (possibly because its message has already been transmitted successfully). We note that if there are η arrival points per slot then taking the limit as $p \rightarrow 0$, $\eta \rightarrow \infty$ while preserving the product $S \triangleq \eta p$, the arrival process becomes Poisson with parameter S .

Capetanakis [5], [6] described the following multiple access protocol for Poisson arrivals. Access to the channel is only permitted at the start of a service "epoch." All messages generated during one epoch are transmitted in the next epoch. Whenever a collision occurs, each station involved in that collision tosses a fair binary coin. Those stations tossing 1's retransmit their messages immediately; those tossing 0's are forbidden from participating in the collision resolution algorithm until all stations that just tossed 1's have successfully transmitted their messages. The capacity of this basic protocol is ≈ 0.34 , but many refine-

ments have subsequently been suggested. Capetanakis modified this basic algorithm to increase the maximum throughput to ≈ 0.43 . Gallager [7] and Tsybakov and Mikhailov [8] independently devised a first-come first-serve (FCFS) protocol with a capacity of ≈ 0.487 . Humblet and Mosely [9] did some fine tuning on the Gallager-Tsybakov algorithm to achieve a maximum throughput of ≈ 0.488 with a FCFS protocol.

In general it is not obvious how an optimal infinite population multiple access protocol should operate. Complicated sequences of intersecting sets could be enabled to gather information about the distribution of messages. Without describing any specific optimal protocols, Pippenger [1] has proven an upper bound ξ_p on capacity. Given any set of arrival points, a protocol enables subsets of these points according to a ternary decision tree. A path through the decision tree terminates when all messages from that set are known to have been transmitted successfully. Each terminal node determines a unique partition of the set of arrival points, with each element containing exactly one message. The bound is obtained in the limit as the expected number of messages in the set of arrival points grows to infinity by equating a lower bound on the entropy of valid partitions (i.e., those that could separate messages correctly) with an upper bound on the entropy of the identity of the terminal node, given that a fraction ξ_p of the time the protocol must take the "success" branch in the decision tree. In the Poisson limit as $p \rightarrow 0$, Pippenger's bound on capacity is ≈ 0.744 .

In this paper we take a more direct approach. We show that optimal protocols can be found for systems in which a helpful "genie" provides certain extra information at no cost to the protocol. The performance of optimal *genie-aided* protocols must be an upper bound on the performance of optimal *unaided* protocols, since the optimal genie-aided protocol can always simulate an unaided protocol by ignoring the genie's information. However, this upper bound need not be attainable.

As a trivial example, perfect utilization of the channel would be possible if the genie were to examine all the arrival points and publicly label each point to be idle or busy. Thus the key to this approach is in the selection of some particular information that does not make the collision resolution problem "too easy" but still allows one to make meaningful statements about the performance of an optimal protocol.

Below, we shall examine protocols that are aided by a genie who publicly labels, at no cost, two busy points (and possibly some idle points) from each collision. We show that the labeling can be done in such a way that all unlabeled points remain effectively unexamined (i.e., the property that the points are independently busy with the probability p is preserved), so that we can get no further information from a collision with genie labeling. This allows us to find the *exact* capacity for all $p \gtrsim 0.568$, and tighten Pippenger's upper bound for Poisson arrivals to ≈ 0.6731 .

II. OPTIMAL GENIE-AIDED ALGORITHMS FOR COLLISION RESOLUTION

Let us assume that the genie uses the following algorithm to label points from each collision. If no previously known busy points were enabled, the genie examines the enabled points chronologically and publicly labels the *first two* busy points that were enabled (thus implicitly labeling every other examined point to be idle). If one previously known busy point was enabled, the genie first labels the known busy point (giving the protocol no new information) and then publicly labels the *first* previously unknown busy point that was enabled (also possibly creating some known idle points). If at least two known busy points were enabled—a certain collision—the genie trivially labels any two known busy points, giving no new information. (This last case clearly cannot be optimal and need not be considered below.)

We note that it is not necessary for the genie to examine the enabled set in a first-come first-served order. The genie may examine the arrival points arbitrarily as long as *every* point that is examined is labeled to be idle or busy. Since the Bernoulli trials were initially independent, all unlabeled arrival points remain effectively unexamined with probability p of being busy even with perfect information about the labeled points. Thus, any genie-aided collision resolution algorithm faces only known busy points, known idle points and unexamined arrival points. It remains to find an optimal genie-aided collision resolution algorithm and to determine the capacity of a protocol that uses that algorithm.

We now show that without loss of generality the search for an optimal algorithm can be restricted to those algorithms that separately enable each known busy point. Let A be any genie-aided algorithm that sometimes enables both a single known busy point and some unexamined arrival points. Define a new algorithm A' that simulates the behavior of A but makes the following modification. Whenever A would enable both one known busy point and $k > 0$ unexamined arrival points—giving a success with probability $(1-p)^k$ and a collision (from which the genie labels *one* new busy point) with probability $1 - (1-p)^k$ — A' enables either one known busy point (and no unexamined arrival points) or all remaining unexamined arrival points (and no known busy points) with probabilities $(1-p)^k$ and $1 - (1-p)^k$, respectively. Should it choose the former, A' resumes its simulation of A as if a success had occurred. Should it choose the latter, there is certain to be a collision from which the genie labels *two* new busy points; A' interrupts its simulation for one slot to transmit successfully one such point and thereafter resumes its simulation of A as if a collision had occurred. There is thus *perfect* utilization of the channel over all slots for which A' interrupts its simulation of A . In addition, the simulation is a faithful probabilistic replica of A : it achieves a successful message transmission with probability $(1-p)^k$ (maintaining the same throughput as A), and it either increases the number of known busy points by one with probability

$1 - (1-p)^k$, or decreases the number of known busy points by one with probability $(1-p)^k$. Thus A' must have at least as high a throughput as A . It follows that no genie-aided algorithm can have a higher capacity than the best genie-aided algorithm that chooses each enabled set to be either a single known busy point or a set of unexamined arrival points. Without loss of efficiency, this may clearly be done FCFS.

For any such genie-aided algorithm, a new period of activity begins whenever the algorithm enables some (possibly random) number N of unexamined arrival points. Each idle period lasts for one slot and transmits successfully no messages, each success period lasts for one slot and transmits successfully one message, and each collision period lasts for three slots and transmits successfully the two genie labelled messages. Over all periods where a particular value of N is chosen, the conditional genie-aided throughput, $\tilde{\rho}_N$, is found from a renewal argument to be

$$\begin{aligned}\tilde{\rho}_N &\triangleq \frac{S_N + 2C_N}{I_N + S_N + 3C_N} = \frac{2 - 2I_N - S_N}{3 - 2I_N - 2S_N} \\ &= 1 - \frac{1 - S_N}{3 - 2I_N - 2S_N},\end{aligned}\quad (1)$$

where $I_N = (1-p)^N$, $S_N = Np(1-p)^{N-1}$, and $C_N = 1 - I_N - S_N$ are the probabilities that enabling N Bernoulli arrival points gives an idle slot, a success or a collision, respectively. Since the unconditional throughput is a convex combination of $\{\tilde{\rho}_N\}$, it cannot exceed $\tilde{\rho}_{N^*}$, where N^* achieves maximum conditional throughput.

We have thus established that an optimal genie-aided strategy is to enable the first known busy point if there is one, or to enable some fixed number N^* of unexamined arrival points otherwise. This optimal strategy also transmits all messages in a first-come first-served order. It remains to determine N^* for all p .

Fact: Let f be a function defined on the positive integers. If either $f_k \geq f_{k+1}$ holds whenever $f_{k-1} \geq f_k$, or $f_k \geq f_{k-1}$ holds whenever $f_{k+1} \geq f_k$, then f is unimodal.

Proof: Let f_i, f_j be two distinct strict local maxima and, without loss of generality, let $j > i$. Then, by the first condition above, we must have

$$f_i \geq f_{i+1} \geq \dots \geq f_j,$$

which contradicts that f_j is a strict local maximum. Similarly, by the second condition, we must have

$$f_j \geq f_{j-1} \geq \dots \geq f_i,$$

which contradicts that f_i is a strict local maximum. Q.E.D.

Theorem 1: For any fixed Bernoulli probability p , sequence $\{\tilde{\rho}_N\}$ is unimodal in N .

Proof: Let $\tilde{\rho}_N \geq \tilde{\rho}_{N+1}$. Then

$$1 - \frac{1 - S_N}{3 - 2I_N - 2S_N} \geq 1 - \frac{1 - S_{N+1}}{3 - 2I_{N+1} - 2S_{N+1}},$$

or

$$S_N - 2S_N I_{N+1} - 2I_N \geq S_{N+1} - 2S_{N+1} I_N - 2I_{N+1}. \quad (2)$$

But $I_{N+1} = (1 - p)I_N$, and $S_{N+1} = pI_N + (1 - p)S_N$, so that (2) gives

$$S_N \geq I_N(3 - 2I_N),$$

or

$$\frac{Np}{1-p} \geq 3 - 2I_N. \quad (3)$$

Hence, to show that $\tilde{p}_{N+1} \geq \tilde{p}_{N+2}$, it suffices to show

$$\frac{(N+1)p}{1-p} \geq 3 - 2I_{N+1} = 3 - 2I_N + 2pI_N.$$

Because of (3), this inequality holds if

$$\frac{p}{1-p} \geq 2pI_N$$

or

$$\frac{1}{2} \geq (1-p)^{N+1}, \quad (4)$$

which is clearly true for $p \geq \frac{1}{2}$. We thus assume $p < \frac{1}{2}$. Since $I_N \leq 1$, it must be the case from (3) that $N \geq (1/p) - 1$. Thus to show (4), it suffices to show that

$$\frac{1}{2} \geq (1-p) \cdot (1-p)^{(1/p)-1}. \quad (5)$$

Let $v \triangleq 1/p > 2$. Since it is well-known that $(1 - 1/v)^{v-1}$ decreases monotonically from 1 to $1/e$ as v increases from 1 to ∞ , (5) follows for all v if it holds for the smallest value, namely $v = 2$, where it is clearly true. Q.E.D.

Theorem 2: For fixed N , there is a unique solution, p_N , to $\tilde{p}_N(p) = \tilde{p}_{N+1}(p)$ for p in the range $0 < p < 1$. If $p < p_N$, then $\tilde{p}_N(p) < \tilde{p}_{N+1}(p)$; if $p > p_N$, then $\tilde{p}_N(p) > \tilde{p}_{N+1}(p)$.

Proof: If $\tilde{p}_N(p) = \tilde{p}_{N+1}(p)$, then (3) must be true as an equality. Let $f(p) \triangleq Np/(1-p)$ and $g(p) \triangleq 3 - 2(1-p)^N$ be the left and right sides of (3), respectively. We now show that assuming the existence of two solutions, $f(p_1) = g(p_1)$ and $f(p_2) = g(p_2)$ for $0 < p_1 < p_2 < 1$, leads to a contradiction. Since $f''(p) = 2N/(1-p)^3$ and $g''(p) = -2N(N-1)(1-p)^{N-2}$, $f(p)$ is strictly convex while $g(p)$ is strictly concave for $0 < p < 1$. For any p_0 such that $0 < p_0 < p_1$, choose α to satisfy $p_1 = \alpha p_0 + (1-\alpha)p_2$. Then by convexity

$$f(p_1) < \alpha f(p_0) + (1-\alpha)f(p_2).$$

Similarly by concavity

$$g(p_1) > \alpha g(p_0) + (1-\alpha)g(p_2).$$

Thus

$$f(p_0) > g(p_0)$$

for all p_0 . Since f and g are continuous and differentiable for $0 \leq p < 1$ and $f(p) \rightarrow 0$ and $g(p) \rightarrow 1$ in the limit as $p \rightarrow 0$, we have an obvious contradiction. There can thus

be at most one solution, p_N . Such a solution must exist, however, since $f(p) \rightarrow \infty$ and $g(p) \rightarrow 3$ as $p \rightarrow 1$. Q.E.D.

Corollary 1: N^* , the optimal number of unexamined arrival points to enable simultaneously, is a nonincreasing function of p .

Proof: Choose any p and determine $N^*(p)$. By Theorem 1, $\tilde{p}_{N^*(p)} \geq \tilde{p}_{N^*(p)+1}$. By Theorem 2, since equality can occur for only one value of p , $\tilde{p}_{N^*(p')} > \tilde{p}_{N^*(p)+1}(p')$ must hold for all $p' > p$. But \tilde{p}_N is unimodal in N , so $N^*(p') \leq N^*(p)$. Q.E.D.

Corollary 2: For $p > 0$, N^* decreases in unit steps as p increases.

Proof: It is sufficient to show that $p_N \leq p_{N-1}$. We thus assume $p_N > p_{N-1}$ and consider $\hat{p} \in (p_{N-1}, p_N)$. Then by Theorem 2 we must have $\tilde{p}_{N+1}(\hat{p}) > \tilde{p}_N(\hat{p})$ since $\hat{p} < p_N$, and $\tilde{p}_{N-1}(\hat{p}) \leq \tilde{p}_N(\hat{p})$ since $\hat{p} > p_{N-1}$, which contradicts that \tilde{p} is unimodal in N . Q.E.D.

We have thus established that \tilde{p}_N is an upper bound to the capacity of optimal protocols for all p in the interval $[p_N, p_{N-1})$.

III. CALCULATING THE BOUNDS

Before computing some specific upper and lower bounds on the capacity of optimal protocols, we wish to establish that capacity (of optimal protocols) must be a *nondecreasing* function of p . This is true because we can always simulate an arrival sequence with Bernoulli probability p_1 given an arrival sequence with $p_2 > p_1$. The next arrival point in this simulated arrival sequence is independently defined either to be the next arrival point from the real arrival sequence (with probability p_1/p_2) or to be empty (with probability $1 - p_1/p_2$). Thus the performance of any protocol for p_1 can be achieved given p_2 by applying the protocol to the simulated arrival sequence.

We note in particular that if $p_1 > 0$ then the *randomized binomial* strategy for p_2 that selects k arrival points with probability $\binom{N}{k} (p_1/p_2)^k (1 - p_1/p_2)^{N-k}$, $0 \leq k \leq N$, enables exactly n busy points with probability

$$\sum_{k=n}^N \binom{N}{k} (p_1/p_2)^k (1 - p_1/p_2)^{N-k} \binom{k}{n} p_2^n (1 - p_2)^{k-n} = \binom{N}{n} p_1^n (1 - p_1)^{N-n}.$$

This distribution is identical to the probability of enabling exactly n busy points by the *fixed* strategy for p_1 that selects N arrival points. Similarly in the Poisson case where $p_1 = 0$, the randomized Poisson strategy for p_2 that selects $k \geq 0$ arrival points with probability

$$\frac{(\lambda/p_2)^k}{k!} e^{-\lambda/p_2}$$

achieves the same distribution for the number of enabled busy points as the fixed strategy for Poisson arrivals that

enables a set with parameter λ . It follows that the performance of the previously described optimal genie-aided protocol for p_1 can be attained given $p_2 > p_1$ by defining a randomized genie-aided protocol. However, since the throughput will be a convex combination of the throughputs of several fixed strategies, there can be no advantage in randomization.

We now consider a Bernoulli arrival process with $p > 1/\sqrt{2} \approx 0.7071$. We note from (1) that $\bar{\rho}_2 = 2p/(1+2p^2) < p$ for all $p > 1/\sqrt{2}$. But since $\bar{\rho}_1 = p$ and $\bar{\rho}_N$ is unimodal in N , no genie-aided protocol (and hence no unaided protocol) can achieve a throughput exceeding p for any $p > 1/\sqrt{2}$. However, enabling individual arrival points whenever there is a backlog of at least one point (i.e., "TDMA") achieves a capacity of p , does not require the genie's help, and is thus feasible and optimal. We note that $1/\sqrt{2}$ must also be an upper bound for all $p \leq 1/\sqrt{2}$ since we have shown that capacity is a nondecreasing function of p .

Similarly using *pairwise* enabling (i.e., $N = 2$) whenever there is a backlog of at least two points is optimal for p between $1/\sqrt{2}$ and the solution of

$$\bar{\rho}_3 = \frac{3p - p^3}{1 + 6p^2 - 4p^3} = \frac{2p}{1 + 2p^2} = \bar{\rho}_2,$$

namely $p \approx 0.568$, where $\bar{\rho}_2 = \bar{\rho}_3 \approx 0.6904$. We can clearly label two busy points from a collision if only two arrival points were enabled, so pairwise enabling is feasible without the genie's help, and thus this is an optimal protocol in the range $0.568 \lesssim p \leq 1/\sqrt{2}$. We may continue to numerically evaluate the boundary where $N \geq 3$, but the optimal protocols now do require the genie's information. In the Poisson limit where $p \rightarrow 0$, N becomes infinite and thus the maximum genie-aided throughput is

$$\bar{\rho}_\infty = \frac{2 - (2 + \lambda)e^{-\lambda}}{3 - 2(1 + \lambda)e^{-\lambda}}.$$

This maximum occurs at $\lambda \approx 2.89$, the solution of $3 - \lambda = 2e^{-\lambda}$. Hence for the case of Poisson arrivals, the throughput cannot exceed $\bar{\rho}_\infty \approx 0.6731$. This completes the calculation of an upper bound on capacity for all p .

We now find a lower bound on capacity in the region $0.568 \gtrsim p > 0$, where our optimal genie-aided protocols simultaneously enable more than two arrival points. Since we cannot label two busy points without the genie's help if a collision occurs when $N \geq 3$, the performance of an optimal genie-aided protocol need not necessarily be attainable without the genie's help. However the performance of any feasible protocol does form a lower bound on capacity. We have also shown that capacity is a nondecreasing function of p , so the capacity of Mosely's algorithm for Poisson arrivals [9], ≈ 0.488 , can be used as a lower bound for all p . We thus obtain a lower bound by taking the maximum of the capacity of Mosely's algorithm and the capacities as a function of p of TDMA and pairwise enabling (the $N = 1$ and $N = 2$ cases, respectively).




Fig. 1. Throughput versus Bernoulli probability.

The feasible region can be expanded upwards slightly by considering contention among more than two Bernoulli trials. We consider three obvious FCFS collision resolution algorithms for $N = 3$. For the first case, assume that when a collision occurs, individual arrival points from the collision are enabled until two messages are transmitted. (There can be no advantage in separately enabling the third arrival point if the first two points were busy because it is effectively unexamined.) Two extra slots resolve the collision if the first two arrival points are busy. Otherwise three extra slots are required. Hence, for this first algorithm, the throughput is given by

$$\frac{3p - p^3}{1 + 2[p^3 + p^2(1 - p)] + 3[2p^2(1 - p)]} = \frac{3p - p^3}{1 + 8p^2 - 6p^3}. \quad (6)$$

In the second case, we immediately enable a *pair* of arrival points whenever a collision occurs. There will either be a success if exactly one of the first two arrival points is idle (and we are done after enabling the third arrival point, now known to be busy), or a further collision if both are busy. Here again, the collision is resolved after exactly two messages have been transmitted, giving a throughput of

$$\frac{3p - p^3}{1 + 3[p^3 + p^2(1 - p)] + 2[2p^2(1 - p)]} = \frac{3p - p^3}{1 + 7p^2 - 4p^3}. \quad (7)$$

Finally, we may first enable a single arrival point from a collision. If that arrival point is idle, the remaining two points (now known to be busy) are separately enabled. If it

is busy, both remaining arrival points are enabled, possibly resulting in a further collision. In this case, collision resolution takes two extra slots if either of the last two arrival points is empty, three extra if the first point is empty, and four extra if all arrival points contained messages. All messages in the enabled set are always transmitted, and the throughput is

$$\frac{3p}{1 + 2[2p^2(1 - p)] + 3[p^2(1 - p)] + 4[p^3]} = \frac{3p}{1 + 7p^2 - 3p^3}. \quad (8)$$

As one can easily verify, (8) exceeds both (6) and (7) in the range $0.206 \lesssim p \lesssim 0.430$, where (8) exceeds both the throughput of pairwise enabling and the lower bound obtained from Mosely's algorithm.

Fig. 1 plots throughput ρ against the probability p of a message arriving at an arrival point. Pippenger's upper bound is shown. In addition the new upper bound presented above is shown delimiting the unattainable region. Mosely's FCFS algorithm, the above FCFS algorithm for $N = 3$, pairwise enabling (the $N = 2$ case described above), and TDMA on arrival points (the $N = 1$ case described above) delimit the attainable region.

IV. CONCLUSION

For all Bernoulli arrival processes with $p \gtrsim 0.568$, optimal algorithms are known: pairwise enabling for $0.568 \lesssim p < 1/\sqrt{2}$ and TDMA for $p \geq 1/\sqrt{2}$. Both algorithms have simple forms and obey FCFS. For smaller p , there is still a range of uncertainty about the maximum attainable throughput.

Even though the new upper bound for Poisson traffic is above $2/3$, we agree with the common conjecture that $1/2$ is the true capacity. However, sharpening the present upper bound from ≈ 0.6731 will seemingly require a more complex argument.

ACKNOWLEDGMENT

In appreciation for their inspiration and many helpful suggestions, the author would like to thank Professors James L. Massey and Leonard Kleinrock.

REFERENCES

- [1] N. Pippenger, "Bounds on the performance of protocols for a multiple-access broadcast channel," *IEEE Trans. Inform. Theory*, IT-27(2), pp. 145-151, March, 1981. Preprint: Res. Rep. RC 7742 (# 33525), Mathematical Sciences Department, IBM Thomas J. Watson Research Center Yorktown Heights, New York, June, 1979.
- [2] B. S. Tsybakov and V. A. Mikhailov, "Ergodicity of a slotted aloha system," *Problems of Information Transmission*, Plenum Publishing Corp., 1980, (translated from Russian).
- [3] N. Abramson, "Packet switching with satellites," *AFIPS Conf. Proc.*, NCC 42, pp. 695-703, June, 1973.
- [4] J. Little, "A proof of the queueing formula $L = \lambda W$," *Operations Research*, 9(2), pp. 383-387, March, 1961.
- [5] J. I. Capetanakis, "Tree algorithms for packet broadcast channels," *IEEE Trans. Inform. Theory* IT-25, pp. 505-515, September, 1979.
- [6] J. I. Capetanakis, "Generalized TDMA: The multi-accessing tree protocol," *IEEE Trans. Commun.* COM-27, pp. 1476-1484, Oct., 1979.
- [7] R. G. Gallager, "Conflict resolution in random access broadcast networks," *Proc. AFOSR Workshop Commun. Theory and Appl.*, pp. 74-76, Sept. 17-20, 1978.
- [8] B. S. Tsybakov, M. A. Berkovskii, N. D. Vvedenskaja, V. A. Mikhailov, and S. P. Fedorov, "Methods of random multiple access," *Fifth International Symposium Inform. Theory*, July 7-9, 1979.
- [9] J. Mosely, "An efficient contention resolution algorithm for multiple access channels," LIDS-TH-918, Laboratory for Information and Decision Systems, MIT, Cambridge, MA, June, 1979.