SIAM J. on Computing
Vol. 0, No. 1, pp. 000–000, October 1995

000

ANALYSIS OF BACKOFF PROTOCOLS FOR MULTIPLE ACCESS CHANNELS *

JOHAN HÅSTAD† AND TOM LEIGHTON ‡ AND BRIAN ROGOFF‡

Abstract. In this paper, we analyze the stochastic behavior of backoff protocols for multiple access channels such as the Ethernet. In particular, we prove that binary exponential backoff is unstable if the arrival rate of new messages at each station is λ/N for any $\lambda > \frac{1}{2}$ and the number of stations N, is sufficiently large. For small N we prove that $\lambda \geq \lambda_0 + \frac{1}{4N-2}$ implies instability where $\lambda_0 \approx .567$. More importantly, we also prove that any superlinear polynomial backoff protocol (e.g., quadratic backoff) is stable for any set of arrival rates that sum to less than one, and any number of stations. The results significantly extend the previous work in the area, and provide the first examples of acknowledgment based protocols known to be stable for a nonnegligible overall arrival rate distributed over an arbitrarily large number of stations. The results also disprove a popular assumption that exponential backoff is the best choice among acknowledgment based protocols for systems with large overall arrival rates. Finally, we prove that any linear or sublinear backoff protocol is unstable if the arrival rate at each station is $\frac{\lambda}{N}$ for any fixed λ and sufficiently large N.

Key words. ethernet, backoff protocols, markov chains, stochastic stability

AMS subject classifications. 60J10, 68M10, 90B12

1. Introduction.

Multiple access channels provide a simple and efficient means of communication in distributed systems. A typical example is the Ethernet [7], a local area network where the channel consists of a tree made out of coaxial cable. When a station wants to send a message to one or more stations on the Ethernet, the sending station simply broadcasts the message throughout the entire system. Everyone, including the intended stations, then receives the message provided that there was no interference from other stations trying to send messages at the same time.

In order to reduce the chance of interference, stations check to make sure that the channel is clear before attempting to transmit a message. At first glance, it might seem that this precaution eliminates the possibility of a collision since the probability that two stations try to send at exactly the same instant in time is virtually zero. Unfortunately, collisions can still occur, since there is a nonnegligible delay between the time when a station begins to transmit and the other stations detect the transmission. Hence, if two or more stations attempt to transmit within this window of time, a collision will occur.

*Received by the editors Month ?, 199?; accepted by teh editors Month, x, xxxx. This research was supported by Air Force Contract AFOSR-86-0078, Army Contract DAAL-03-86-K-0171, an NSF Presidential Young Investigator Award with matching funds from IBM, Xerox and AT&T Bell Labs, and an IBM Postdoctoral Fellowship..

†Department of Numerical Analysis and Computing Science, Royal Institute of Technology (KTH), S-100 44 Stockholm, SWEDEN, (johanh@nada.kth.se).

‡Mathematics Department and Laboratory of Computer Science, Massachusetts Institute of Technology,Cambridge, MA 02139 (ftl@math.mit.edu).

In the case of a collision, none of the messages is sent. Instead, the collision is detected and the messages are queued at their respective stations for retransmission at some point in the future. Of course, it would not make sense to retransmit right away since this would immediately result in another collision. Rather, packets are retransmitted according to a protocol that is often probabilistic in nature. For example, messages in an Ethernet are retransmitted again after T steps where T is selected randomly from $\{1, 2, 3, ..., 2^{\min(10,b)}\}$ and b is the number of times the station has tried to send the packet but failed. This is one of a class of protocols generally referred to as *exponential backoff*.

The success of a protocol can be measured in several ways. For example, we might be interested in the average waiting time W_{ave} incurred by a message before it is successfully transmitted. Alternatively, we might consider the average number of waiting messages over time L_{ave} to be a better measure. Actually, these measurements are closely related. In fact, $L_{ave} = \lambda W_{ave}$ with probability one where λ is the overall arrival rate of messages into the system over time [12].

For a protocol to be useful, it is crucial that $Ex[L_{ave}]$ and $Ex[W_{ave}]$ be small. In particular, we will want $Ex[L_{ave}]$ and $Ex[W_{ave}]$ to be finite. Note that this is a stronger condition than insisting only that L_{ave} and W_{ave} be finite with probability one. For example, consider the situation when $W_{ave} = 2^i$ with probability 2^{-i} for $i = 1, 2, 3 \ldots$.

Another measure of system performance that is often of interest to statisticians is the expected time $Ex[T_{ret}]$ for the system to return to the *start state* (i.e. the state where all queues are empty). Of course, we will want this time to be as small as possible, and, in particular, we will want it to be finite.

Protocols for which $Ex[L_{ave}]$, $Ex[W_{ave}]$ and $Ex[T_{ret}]$ are finite are said to be *stable*. Protocols for which all the measures diverge are said to be it unstable. Note that it is conceivable that there are protocols that are neither stable nor unstable as we have defined the terms here, since it might be that case that $Ex[T_{ret}]$ is finite but $Ex[W_{ave}]$ diverges for some protocol. However, all of the protocols considered in this paper are shown to be either entirely stable or entirely unstable in the sense defined above. In fact, the only reason we use these somewhat nonstandard definitions is that we want to encompass as many of the conflicting definitions of stability and instability in the literature as possible with our methods.

The *throughput rate* (i.e., the average rate of successful transmissions) is not a dominant concern. This is because the throughput rate is guaranteed to equal the arrival rate with probability one if the protocol is stable, but not vice-versa. As an example, consider a 1-station system in which the sole station transmits with probability one if it has a message and in which the station receives a pair of new messages at each step with probability 1/2. It is not difficult to show that with probability one, both the arrival and throughput rates for this system are 1 but that L_{ave} and W_{ave} diverge over time.

The development and analysis of transmission protocols that minimize average waiting time has been the subject of a great deal of work [1– 7,9–17]. We summarize some of this work in Section 2. Of greatest concern in this paper is the work on acknowledgment based protocols. An *acknowledgment based protocol* is one for which each station's transmission protocol is based only on its own history of successes and failures. In particular, the station is not assumed to have any knowledge of other stations' successes or failures or even of the number of stations in the system N.

Our present work is focused on a subset of acknowledgment based protocols known as backoff protocols. A *backoff protocol* is one for which each station i containing a message transmits with probability $f(b_i)$ where f is a predetermined function and b_i (the *backoff counter* at station i) is the number of past consecutive failures by station i. After each successful transmission, b_i is reset to zero. After each failure, b_i is augmented by one. The value of b_i is left unchanged if no transmission is attempted.

Previous work has mostly centered on *binary exponential backoff* (for which $f(b) = 2^{-b}$), although other schemes such as *linear backoff* (for which $f(b) = \frac{1}{b+1}$) and *constant backoff* (for which $f(b)$ is simply a constant) have also been considered. Unfortunately, most of this work has been experimental and/or has depended on simplifying assumptions (e.g., that N is infinite) that render the consequences of any analysis less meaningful. Exceptions include some work on constant backoff (which serves as the basis for the Aloha protocol, but which is inherently unstable for fixed backoff and sufficiently large N) [13, 17], and the work of Goodman, Greenberg, Madras and March [2], who proved that for any N, there is a λ' for which exponential backoff has finite $Ex[T_{ret}]$ provided that the arrivals at station i are Bernoulli distributed with mean λ_i for $1 \leq i \leq N$ where $\lambda = \sum_{i=1}^{N} \lambda_i \leq \lambda'$. Unfortunately, λ' tends to zero as N increases and the question concerning stability for nonvanishing λ and large N remained open. In the case when $N = 2$, Goodman et.al. also proved that exponential backoff has finite $Ex[T_{ret}]$ if λ_1 and λ_2 are at most 0.15, and infinite $Ex[T_{ret}]$ if $\lambda_1 > 1/2$ and $\lambda_2 > 0$.

In this paper, we redirect the focus from exponential backoff to other protocols. Among other things we show that exponential backoff is unstable whenever $\lambda_i \geq \frac{\lambda}{N}$ for $1 \leq i \leq N$ and $\lambda > 0.567 + \frac{1}{4N-2}$, or when $\lambda > \frac{1}{2}$ and N is sufficiently large. The result is not very surprising given the existing experimental data, but it does establish formal limits on the usefulness of exponential backoff. We also prove a much stronger and more important result concerning the stability of polynomial backoff protocols. In particular, we prove that if the arrivals at station i are Bernoulli distributed with rate λ_i then $f(b) = (b+1)^{-\alpha}$ backoff is stable for any constant $\alpha > 1$, any N and any $\{\lambda_i \mid \lambda = \sum_{i=1}^{N} \lambda_i < 1\}$. In terms of stability, the result is the strongest possible since any protocol is unstable if the overall arrival rate λ is one or larger. The result also provides the first example of an acknowledgment based protocol known to be stable for nonvanishing λ and large N, and proves that polynomial backoff protocols are superior to exponential backoff when λ is large.

The constraint that α be greater than one is crucial to the stability of polynomial backoff. In fact, we also prove that for any $\alpha \leq 1$, $f(b) = (b+1)^{-\alpha}$ backoff is unstable for any evenly distributed arrival rate λ and sufficiently large N.

Once a protocol has been found to be stable, the next step is to determine the precise values of $Ex[L_{ave}]$ and $Ex[W_{ave}]$. In particular, it is interesting to analyze the dependence of $Ex[L_{ave}]$ and $Ex[W_{ave}]$ on λ and N. Unfortunately, our current best estimates for these values are fairly weak. Whereas we do prove that $Ex[L_{ave}]$ and $Ex[W_{ave}]$ must grow polynomially as a function of N for almost any backoff protocol, we only upper bound this growth by an exponential function of N.

Quantifying the nature of an instability can also be of interest, particularly if the protocol is to be implemented in practice. In the cases of exponential backoff with high arrival rates and linear or sublinear backoff with large numbers of stations, we show that $Ex[L_{ave}]$ grows linearly over time, the worst possible scenario.

As a crucial aid in guiding our research, we performed computer simulations of several backoff protocols for various arrival rates and numbers of stations. Some of the data obtained can be found in Section 7. This data does suggest that quadratic backoff is a very competitive algorithm in practice. Linear backoff seems better if the number of stations is small or the load is minimal. Exponential backoff seems better only when the queues are massive.

The remainder of the paper is divided as follows. In Section 2, we describe the models for communication protocols more formally, introduce some further notation, and comment on the relevance of past research to our current work. In Section 3, we prove that polynomial backoff is stable for any arrival rate less than 1. In Section 4, we examine the dependence of L_{ave} and W_{ave} on λ and N. The instability of exponential backoff protocols is established in Section 5. In Section 6, we show that linear and sublinear backoff protocols are unstable for any fixed arrival rate and sufficiently large N. Section 7 contains some experimental data, and we conclude with some remarks and topics for research in Section 8.

2. Preliminaries.

2.1. Our model.

In this paper, we follow the model of backoff protocols adopted in [2]. In this model, time is partitioned into equal length intervals called *steps*. At the beginning of each step, a new message arrives at station i with probability λ_i for $1 \leq i \leq N$ where N is the number of stations in the system. The arrival of new messages is assumed to be independent over time and among stations. The *overall arrival rate* is defined to be $\lambda = \sum_{i=1}^{N} \lambda_i$. Arriving messages are added to the end of the *queues* located at each station. No limit is placed on the size of the queues, and if the system is unstable, they could become arbitrarily long over time.

The backoff protocol is governed by a function $f(x)$ defined in advance. At each step, the ith station attempts a transmission if it has a nonempty queue (allowing for the arrival of a new packet at the beginning of the step) with probability $f(b_i)$ where b_i is the value of the *backoff counter* at station i. For any set of b_i's, these probabilities are assumed to be independent over time and among stations. The backoff counters are initially zero. The ith counter is augmented by one whenever the ith station attempts to transmit but fails due to a collision. The ith counter is reset to zero whenever the ith station transmits successfully. If the ith station does not attempt to transmit, then the backoff counter is not changed. Confirmation of a collision or a successful transmission takes place during the same step in which the transmission was attempted. In addition, the message lengths are assumed to be smaller than the duration of a step.

The two important measures of efficiency are the average number of messages queued in the system at the end of each step L_{ave} and the average number of steps that each message must wait before it is sent W_{ave}. Since $L_{ave} = \lambda W_{ave}$ with probability 1, we will henceforth express our results in terms of L_{ave}.

2.2. Relevance of the model to reality.

Our mathematical model differs from reality (e.g., the Ethernet) in several respects. We summarize these differences and their significance in the following paragraphs.

Upper bound on backoff counter. In the Ethernet, the backoff counter is never allowed to exceed a specified value b_{max}, whereas in the mathematical model it is allowed to become arbitrarily large. One mathematical problem with placing an upper bound on the backoff counter is that any such protocol becomes unstable for any fixed λ and large enough N. The reason is that for very large N, the system will eventually reach a state where almost every station has many messages queued. Once this happens, then with nonzero probability, the channel will become dominated by collisions and the throughput will be forever reduced to a trickle. The situation is less clear if the bound b_{max} is allowed to depend on N, but then individual stations would need to be informed about the number of other stations in the system. Of course, it might well be that reasonable upper bounds on N could be assumed in the computation of a bound b_{max} in practice. Research into such protocols for bounded N might prove to be interesting mathematically as well as useful in practice. For example, see [13, 17] and the references they contain for a discussion of constant backoff protocols.

Termination of undelivered messages. In the Ethernet, messages are discarded if they are not delivered within a specified amount of time. In our model, messages are never discarded and might be held in queue for an arbitrary amount of time. Discarding messages assures stability in the sense that $\mathrm{Ex}[L_{ave}]$ is guaranteed to be finite, but only at the expense of discarding a nonzero fraction of the messages in systems that become unstable if discarding is not allowed.

Distribution of arrivals. In real systems, the distribution of arrivals may not be Bernoulli and may not be independent among stations. However, the Bernoulli and independence assumptions seem as reasonable as any others that are capable of being analyzed. Moreover, our analysis extends to several other natural distributions, and can even be extended to systems where there is dependence among stations. For example, both stability and instability results hold for systems where at most two packets enter the entire system during a single step.

Selection of waiting time before attempted retransmission. In the Ethernet, the ith station attempts to rebroadcast t steps after the last attempt where t is selected uniformly from $\{1, 2, 3, \ldots, 2^{b_i}\}$. In our model, we retransmit at each step with probability 2^{-b_i}. The two methods for computing retransmission times are quite similar, but the former is easier to implement in practice (since it requires fewer random bits) and the latter is easier to analyze (since it is memoryless). It would seem unlikely that the stability results would differ for the two methods, but we have not proved this.

Message length. In reality, message lengths are much longer than the window during which conflicts can arise and be detected. Moreover, message lengths may vary from message to message. In the Ethernet, all messages are restricted to have the same length, and this length is a reasonably large multiple of the window of time used for conflict resolution. This difference between the Ethernet and our model is not as great as it might seem at first, however. The reason is that we can model a system where transmissions are long (but with uniform length) with a system where transmissions have zero length by simply compressing time to squeeze out transmission times altogether. This does not affect the conflict resolution process (which lies at the heart of stability analysis). Rather, we need only adjust the message arrival rates so that there are proportionately fewer arrivals in steps following nontransmissions. Although we will not go through the details here, it is not difficult to extend our stability results to hold for such a modified arrival process. Our instability results are also meaningful in systems with large message lengths, but require some changes. The reason is that we only know how to prove instability in a message-length-M model for arrival rate λ if the corresponding message-length-one model is unstable for arrival rate $\frac{\lambda}{\lambda+(1-\lambda)M}$. Hence, since our instability results for exponential backoff hold only for arrival rates exceeding 0.5, they only imply instability for arrival rates approaching one as M increases. Our linear and sublinear backoff results apply to any constant arrival rates and hence give the same results for any fixed M.

Synchronization. In real systems, time is not partitioned into discrete "windows" because there is no synchronization. In the Ethernet, a station that wants to transmit simply does so when and if the channel is clear locally. Nevertheless, it can be argued that our synchronous model accurately represents an asynchronous system to within a factor of two in window size [6]. Hence this assumption should not have a significant impact on stability analysis.

The bottom line. Whereas our model differs from reality in several notable respects, the differences are not all as important as they first seem. Moreover, the model is the most realistic among those that have been formally studied in the literature. In summary, the real contribution of this work is the development of formal techniques for analyzing communication protocols in multiple access channels, and the observation that protocols such as quadratic backoff may be superior to currently used protocols such as exponential backoff when the number of stations and/or the overall arrival rate is large.

2.3. Other models and results.

Most of the work on protocols for communication in multiple access channels has focused on models for which the number of stations N is infinite [1, 3, 5, 9, 10, 14, 16]. The attraction for study of infinite models is clear. On the one hand, the analysis is simpler since with probability one, no two packets will ever arrive at the same station, and thus the disposition of packets is effectively independent of the station that transmits them (e.g., each message has its own backoff counter and is never contained in a queue with something else). On the other hand, there is the argument that the behavior of a protocol in an infinite model is reflective of its behavior in a real system with a large number of stations.

It has been our experience, however, that infinite model results often have fairly limited relevance to finite systems even when the number of stations is very large. Indeed, the reason is precisely that queuing plays a major role in any finite system (even one with large N), but is nonexistent in the infinite model. As a striking example, we note that Kelly proved in [5] that any polynomial backoff scheme is unstable in the infinite model. More recently, Aldous [1] extended this result to show that exponential backoff is also unstable in the infinite model. Whereas the complete disposition of exponential backoff in a finite model still remains unclear, we show that polynomial backoff is stable for any finite number of stations. Hence, the behavior of backoff protocols in infinite models can be misleading.

The study of infinite models has some implications for our paper, however. For example, the techniques used to prove the instability of backoff protocols in the infinite model can be used to prove that $Ex[L_{ave}]$ grows at least as a polynomial function of N in the finite model. More generally, it appears that a protocol is unstable in an infinite model if and only if the value of $Ex[L_{ave}]$ grows as a function of N in the corresponding finite model. In fact, we follow this strategy in proving lower bounds for $Ex[L_{ave}]$ in Section 4.

There has also been a great deal of work on models that require the use of more information when computing transmission probabilities. Some models use knowledge of the number of stations, or try to approximate the number of stations that wish to transmit by analyzing the past history of the channel activity. Still others try to resolve conflicts among transmitting stations by using a playoff-type system to eventually choose a winner. Such schemes tend to be very stable for input rates up to a fixed threshold (e.g., $1/e$) but unstable for larger input rates. For examples of such models and their analysis see [4, 9, 16] and the references they contain.

As a final note, we point out that many protocols can be made stable for any fixed N and $\lambda < 1$ by simply allowing a transmitting station to empty its queue before allowing anyone else to start. Whereas such an approach may be necessary as a last ditch effort, it is not considered desirable since it allows a single station to dominate the system for a very long time, and since there must be a nontransmission step following the emptying of every queue. The latter constraint is particularly damaging in practice since if the protocol is working well, most queues will be very short, and hence the resulting frequency of nontransmissions is forced to be large (which means the protocol isn't working well after all). In particular, for large λ, we must have $Ex[L_{ave}] \geq \Omega(N)$ for such schemes. Although we prove an even greater asymptotic lower bound on $Ex[L_{ave}]$ for polynomial backoff, backoff protocols perform much better experimentally and are much simpler to implement. In fact, there is some reason to believe that polynomial backoff schemes perform like queue-emptying protocols during the rare times when they get into trouble (indeed, this possible behavior is the basis for our proof of stability), and otherwise behave similarly, but without the need for forced nontransmission steps.

2.4. Markov chains and their analysis.

The performance of most any protocol can be expressed in terms of the behavior of an associated Markov chain. For backoff protocols with a finite number of stations, we associate every possible configuration of backoff counters and queues $(\mathbf{b}, \mathbf{q}) = \{(b_1, \ldots, b_N, q_1, \ldots, q_N) \mid \ b_i \geq 0,\ q_i \geq 0 \ \text{for} \ 1 \leq i \leq N\}$ with a unique state of the Markov chain. The initial state (or origin or zero) is identified with $(0, \ldots, 0)$. The associated infinite Markov chain is time invariant (the transition probabilities do not change with time), irreducible (every state is reachable from every other state) and aperiodic (the probability of being in any state at time t is positive if t is sufficiently large). We will below discuss some properties of Markov chains. For a more detailed discussion we refer to [8].

A Markov chain is said to be *positive recurrent* if the expected time to return to zero $Ex[T_{ret}]$ is finite. It is said to be *transient* if the probability of returning to zero is less than one. Transience is a stronger condition than not being positively recurrent since any transient chain is clearly not positive recurrent, but not vice-versa. For example, an unbiased random walk is neither positive recurrent nor transient. We will also be interested in a third property; namely whether the expected queue size $Ex[L_{ave}]$ over time is finite.

In the literature, a system is often said to be stable if it is positive recurrent. In practical situations, stability more naturally corresponds to the situation when $Ex[L_{ave}]$ is finite. Hence, we adopt a hybrid definition of stability in this paper. In particular, we say that a protocol is *stable* if it satisfies both conditions, and that it is *unstable* if it satisfies neither. Although there are hypothetical examples of systems that satisfy either condition but not the other, the protocols we study either satisfy both or neither. Hence we will be able to classify protocols as stable in the strongest possible sense or unstable in an equally strong sense. Of course, there is also the possibility of the initial state being transient within the domain of instability. In fact, we conjecture that our instability results can be extended to prove the initial states is transient although our techniques do not appear sufficient to prove such an extension.

The predominant method for analyzing the behavior of an infinite Markov chain is by means of a potential (or Lyapanov) function. In our case, a potential function is a map from $(\mathbb{Z}^+)^{2N}$ to $\mathbb{Z}^+$ such that the origin is mapped to zero and the other states are mapped to positive integers. Often, the potential function is directly tied to the measure of concern (e.g., the number of messages held in queues). For example, we will use potential functions of the form

$$c_1 \sum_{i=1}^{N} q_i \pm \sum_{i=1}^{N} f(b_i)^{-c_2} + c_3$$

for some constants $c_1 > 0$, $c_2 \geq 1$, and c_3.

The key step in proving that a protocol is unstable is to find a potential function for which the expected change in potential is at least δ for any state, where δ is a fixed positive constant. This, of course, implies that the expected value of the potential function after t steps is at least δt. By itself, this is not enough to imply instability. For example, consider the simple chain where state i moves to state $\max(i-1, 0)$ with probability $3/4$ and to state $i+1$ with probability $1/4$, and for which the potential of state i is $e^{2i} - 1$. This chain is positive recurrent but the expected change in potential is always at least 1.

If the potential function is natural enough, however, then such an argument can be used to prove instability. For example, we will use a potential function of the form

$$\mathrm{POT}(\mathbf{q}, \mathbf{b}) = C \sum_{i=1}^{N} q_i + \sum_{i=1}^{N} f(b_i)^{-1} - N$$

for exponential backoff. If this potential function grows linearly with time, one can prove that $\mathrm{Ex}[L_{ave}]$ diverges and that the system is unstable. We will prove this formally in Section 5.

The key step in proving that a protocol is stable is to find a potential function for which the expected change in potential is at most $-\delta$ for all but a finite number of states, where δ is a positive constant. Once such a potential function is found for $c_2 \geq 1$, it can then be shown that the associated chain is positive recurrent. To prove that $Ex[L_{ave}]$ is finite one has to study the expected change in POT^2. We will establish this connection in Section 3.

Unfortunately, it is not clear how to find such a potential function for polynomial backoff protocols. In fact, we suspect that there is no such potential function which increases monotonically with the q_i's and the b_i's. Hence, we must follow a somewhat more complicated approach to prove that polynomial backoff is stable.

In particular, we find a potential function for which there is a constant-depth tree of descendent states (not necessarily all of the same depth) emerging from each state for which the expected change in potential computed over these states is at most $-\delta$. In other words, the potential might be expected to increase in the first few steps, but must decrease overall after some larger (but constant) number of steps. Such an argument is still sufficient to prove stability since the performance of such a chain is equivalent (up to constant factors) to the performance of a chain where each tree of descendent states is replaced by direct transitions from the root to the leaves with the appropriate probabilities. The latter chain is then shown to be positive recurrent and stable by the usual approach.

3. Stability of Polynomial Backoff.

In this section, we show that polynomial backoff is stable under the most general assumptions. In particular, we prove the following theorem.

Theorem 3.1: *Let $f(x) = (x+1)^{-\alpha}$ for any $\alpha > 1$. Then, for any number of stations N and any set of arrival probabilities $\lambda_1, \ldots, \lambda_N$ that sum to $\lambda < 1$, the backoff protocol defined by $f(x)$ is stable.*

The overall strategy of the proof is along the lines described in Section 2.4. In particular, we define the potential function

$$POT(\mathbf{q}, \mathbf{b}) = \sum_{i=1}^{N} q_i + \sum_{i=1}^{N} (b_i + 1)^{\alpha+1/2} - N$$

where q_i is the length of the ith queue and b_i is the value of the ith backoff counter and $\mathbf{q}$ and $\mathbf{b}$ are the corresponding vectors. We show that for every state with sufficiently large potential POT, there is a constant-depth tree of descendent states over which the expected decrease in POT^2 is at least δPOT for some fixed constant $\delta > 0$. By a tree of descendant states we mean the following. Starting at a state $\mathbf{q^0}, \mathbf{b^0}$ we follow the system step by step. We observe the system and at each time we decide whether to halt the system or to let it run for another timestep. We always halt the system within a finite number of steps. The total sets of halted states naturally form a tree and the maximal number of steps we observe the system is the depth of the tree. Standard theorems establishing convergence (see [8]) does not seem to apply to this situation and hence to prove Theorem 3.1, we need the following lemma.

Lemma 3.2: *Suppose that there are constants δ, d and V such that for any state $(\mathbf{q}, \mathbf{b})$ which have potential $POT(\mathbf{q}, \mathbf{b}) \geq V$, there is a tree with depth at most d of descendent states over which the expected decrease in POT^2 is at least $\delta POT(\mathbf{q}, \mathbf{b})$. Let T_{retV} denote the time at which the system returns to potential V or less (If $POT(\mathbf{q}, \mathbf{b}) \leq V$ then $T_{retV} = 0$). Then there is another constant c depending only on δ, d and V for which*

$$Ex[\sum_{t=0}^{T_{retV}} L(t)|\ (\mathbf{q^0}, \mathbf{b^0}) = (\mathbf{q}, \mathbf{b})] \leq cPOT^2(\mathbf{q}, \mathbf{b})$$

where $L(t)$ denotes the number of items in the system (i.e. total queue length) at time t.

Proof: We will prove the lemma by induction on time, but we have to be careful since T_{retV} might be infinite, a priori. To overcome the subtleties inherent in dealing with large values of T_{retV}, we define a modified system that is terminated after T steps. In particular, at time T, the system automatically returns to the origin and remains there forever. We then examine

$$E(\mathbf{q}, \mathbf{b}, T) = Ex[\sum_{t=0}^{\min(T, T_{retV})} L(t)|\ (\mathbf{q^0}, \mathbf{b^0}) = (\mathbf{q}, \mathbf{b})]$$

and proceed by induction on time. For $T < 0$ we formally define $E(\mathbf{q}, \mathbf{b}, T) = 0$. Our induction hypothesis is

$$E(\mathbf{q}, \mathbf{b}, T) \leq cPOT^2(\mathbf{q}, \mathbf{b})$$

for all values of $\mathbf{q}, \mathbf{b}$ and T.

Provided that $c \geq 1$, the hypothesis is true for $T = 1$, since $L(0) \leq POT(\mathbf{q}, \mathbf{b})$ and $L(1) = 0$. In addition, the hypothesis is also true if $POT(\mathbf{q}, \mathbf{b}) \leq V$ by definition, since then $T_{retV} = 0$. We next assume that

$$E(\mathbf{q}, \mathbf{b}, T') \leq cPOT^2(\mathbf{q}, \mathbf{b})$$

for all $T' < T$ and any $\mathbf{q}$ and $\mathbf{b}$, and consider the case when the system is terminated at time T.

Let the ith leaf of the tree of descendent states appear with probability p_i, have potential POT_i, and be at depth d_i. Also let L_i denote the sum of $L(t)$ over d_i steps taken to reach the ith leaf. Since at most one item be broadcast at any step, and $L(t) \leq POT(\mathbf{q^t}, \mathbf{b^t})$ we can deduce $L(t-j) \leq POT(\mathbf{q^t}, \mathbf{b^t}) + j$ and hence

$$L_i \leq \sum_{j=0}^{d_i} (POT_i + j) = (d_i + 1)POT_i + \frac{d_i(d_i + 1)}{2}.$$

Thus we can conclude that

$$E(POT, T) \leq \sum_i p_i (L_i + E(POT_i, T - d_i))$$

$$\leq \sum_i p_i((d_i + 1)POT_i + \frac{d_i(d_i + 1)}{2} + cPOT_i^2) \leq$$

$$\leq (d + 1) \sum_i p_i POT_i + \frac{d(d + 1)}{2} + c \sum_i p_i POT_i^2.$$

By the assumption of the lemma we know that

$$\sum_i p_i POT_i^2 \leq POT^2(\mathbf{q}, \mathbf{b}) - \delta POT(\mathbf{q}, \mathbf{b}).$$

A standard convexity argument can be used to show that this implies that

$$\sum_i p_i POT_i \leq POT.$$

Hence,

$$E(\mathbf{q}, \mathbf{b}, T) \leq (d + 1)POT(\mathbf{q}, \mathbf{b}) + \frac{d(d + 1)}{2} + cPOT^2(\mathbf{q}, \mathbf{b}) - c\delta POT(\mathbf{q}, \mathbf{b}).$$

By choosing $c \geq \frac{d+1}{\delta} + \frac{d(d+1)}{2\delta V}$, we can then conclude that $E(\mathbf{q}, \mathbf{b}, T) \leq cPOT^2(\mathbf{q}, \mathbf{b})$, which concludes the induction.

We have now proved that

$$Ex[\sum_{t=0}^{\min(T, T_{retV})} L(t) |\ (\mathbf{q^0}, \mathbf{b^0}) = (\mathbf{q}, \mathbf{b})]] \leq cPOT^2(\mathbf{q}, \mathbf{b})$$

for any T, where the constant c does not depend on T. Assume for the purposes of contradiction that

$$Ex[\sum_{t=0}^{T_{retV}} L(t)|\ (\mathbf{q^0}, \mathbf{b^0}) = (\mathbf{q}, \mathbf{b})]] > cPOT^2(\mathbf{q}, \mathbf{b})$$

for some state $(\mathbf{q}, \mathbf{b})$. Then there would be a finite T for which

$$Ex[\sum_{t=0}^{\min(T,T_{retV})} L(t)|\ (\mathbf{q^0}, \mathbf{b^0}) = (\mathbf{q}, \mathbf{b})]] > cPOT^2(\mathbf{q}, \mathbf{b})$$

which is a contradiction. Hence

$$Ex[\sum_{t=0}^{T_{retV}} L(t)|\ (\mathbf{q^0}, \mathbf{b^0}) = (\mathbf{q}, \mathbf{b})]] \leq cPOT^2(\mathbf{q}, \mathbf{b})$$

as claimed. ♠

Next we have.

Lemma 3.3: *Any system that satisfies the hypothesis of Lemma 3.2, and for which states with potential less than V can only move to states of potential at most $O(V)$ is stable.*

Proof: We use Theorem 14.0.1 of [8]. Let us state this theorem in our vocabulary.

Theorem [14.0.1 from [8]] *Given a Markov chain on a denumerable set which is irreducible and aperiodic and let $f \geq 1$ be a function on its statespace. Then the following conditions are equivalent:*

(i) The chain is positive recurrent with invariant probability measure π and the expected value for f with respect to π is finite.

(ii) There exist a finite set C of states such that

$$\sup_{(\mathbf{q},\mathbf{b})\in C} Ex[\sum_{t=1}^{T_{retC}} f(\mathbf{q^t}, \mathbf{b^t})|\ (\mathbf{q^0}, \mathbf{b^0}) = (\mathbf{q}, \mathbf{b})] < \infty$$

where $T_{retC} > 0$ is the time needed for the chain to return to C after step 0.

Lemma 3.3 follows from this theorem. We use f to be total queue length and C to be the set of states with potential at most V. Then condition (ii) follows from Lemma 3.2 and the conclusion of Lemma 3.3 is then given by (i). ♠

For polynomial backoff protocols with constant N and α, the potential of the system can increase by at most a constant factor at each step. Hence the condition of Lemma 3.3 that states with potential less than V be constrained to move to states with potential $O(V)$ is easily satisfied. Polynomial backoff protocols also satisfy the conditions of Lemma 3.2, although this is much harder to verify. In fact, the bulk of this section will be devoted to establishing the hypothesis of Lemma 3.2. The analysis is divided into four cases, depending on the magnitude of the transmitting probabilities associated with the state. To this end let us define

$$p_i = \begin{cases} \lambda_i & \text{if } b_i = 0 \text{ and } q_i = 0 \\ 1 & \text{if } b_i = 0 \text{ and } q_i > 0 \\ (b_i + 1)^{-\alpha} & \text{if } b_i > 0 \end{cases}$$

to be the probability that the ith station attempts a transmission. The four cases are then

I) $\forall i\ b_i \leq B$,
II) $\exists i\ b_i \geq B$ and $\forall i\ p_i < 1$,
III) $\exists i\ b_i \geq B$, $\exists i\ p_i = 1$, and there exists another i with $p_i \geq \frac{1}{M}$, and
IV) $\exists i\ b_i \geq B$, $\exists i\ p_i = 1$, and for all other i, $p_i \leq \frac{1}{M}$.

The values of B and M are constants to be defined later. Throughout, we will assume that $POT(\mathbf{q}, \mathbf{b}) \geq V$ where V is another large constant to be determined later.

In what follows it will be convenient to let Q_i^+ (Q_i^-) denote the expected increase (decrease) in the potential due to changes in the length of the ith queue. We define B_i^+ and B_i^- in an analogous way, and we let $Q^+ = \sum_{i=1}^{N} Q_i^+$, and define Q^-, B^+ and B^- analogously. We use $Ex[X]$ denote the expected value of random variable X and $\Delta(P)$ to denote the amount change in the quantity P. For example, the expected change in potential will be written as $Ex[\Delta(POT)]$. Lastly, we refer to the ith station as S_i. All these quantities are dependent on the present state $(\mathbf{q}, \mathbf{b})$ but due to readability considerations, we will not make this dependence explicit.

We analyze the cases in order of their difficulty. Case I is by far the most difficult and is saved for last. We start with Case II.

Case II: Without loss of generality we can assume that $\forall i\ p_i < 1, b_1 \geq B$ and $b_i \leq b_1$ for $i > 1$.

We consider a single step of the system and analyze $\Delta(POT^2)$. By definition,

$$Ex[\Delta(POT^2)] = Ex[\Delta(POT^2)|S_1\ succeeds]Pr[S_1\ succeeds]$$

$$+Ex[\Delta(POT^2)|S_1\ does\ not\ succeed]Pr[S_1\ does\ not\ succeed].$$

When S_1 succeeds the potential decreases by at least

$$(b_1+1)^{\alpha+\frac{1}{2}} - N \geq \frac{1}{2}(b_1+1)^{\alpha+\frac{1}{2}}$$

if $B \geq 10N$. This in turn, corresponds to a decrease in magnitude for POT^2 of at least

$$(b_1+1)^{\alpha+\frac{1}{2}}POT - \frac{1}{4}(b_1+1)^{2\alpha+1} \geq \frac{1}{2}(b_1+1)^{\alpha+\frac{1}{2}}POT$$

since $POT \geq \frac{1}{2}(b_1+1)^{\alpha+\frac{1}{2}}$ by definition.

The probability of S_1 succeeding is at least

$$(b_1+1)^{-\alpha}\prod_{i=2}^{N}(1-p_i) \geq (1-\lambda)2^{-N}(b_1+1)^{-\alpha}.$$

Thus the first term in the expression for $Ex[\Delta(POT^2)]$ is at most

$$-\frac{1}{2}(1-\lambda)2^{-N}(b_1+1)^{\frac{1}{2}}POT.$$

To estimate the second term, we use $\Delta Q^+ \leq N$, $\Delta Q^- \geq 0$ and $\Delta B^- \geq 0$ to obtain

$$Ex[\Delta(POT^2)|S_1\ does\ not\ succeed]Pr[S_1\ does\ not\ succeed]$$

$$\leq Ex[\Delta(POT^2)|S_1 \textit{ does not succeed}]$$
$$\leq Ex[(POT + N + \Delta B^+)^2] - POT^2$$
$$\leq N^2 + 2N \cdot POT + 2(POT + N)Ex[\Delta B^+] + Ex[(\Delta B^+)^2]$$

where ΔB^+ is conditioned on the fact that S_1 does not succeed. Note that this does not mean that S_1 tried and failed, since S_1 probably did not even try. In any event, if is not difficult to verify that $Ex(\Delta B^+) \leq N$ and $Ex((\Delta B^+)^2) \leq N(b_1+1)^{\alpha-1} + N^2$. Plugging in and summing, we find that the second term is at most

$$4N^2 + 4N \cdot POT + N(b_1+1)^{\alpha-1}.$$

Combining the two terms with the inequality $POT \geq (b_1+1)^{\alpha-1}$ then gives

$$Ex[\Delta(POT^2)] \quad \leq \quad 4N^2 + 4N \cdot POT + N(b_1+1)^{\alpha-1} - \frac{1}{2}(1-\lambda)2^{-N}(b_1+1)^{\frac{1}{2}}POT \leq -\delta POT$$

provided that $\delta < 1$, $POT \geq N$, and

$$B \geq \frac{c_B N^2 2^{2N}}{(1-\lambda)^2}$$

for some constant c_B independent of N and λ.

Case III: We assume $b_1 \geq B,\ p_2 = 1,\ p_j \geq \frac{1}{M},\ j \neq 2$ and $b_i \leq b_1$ for $i > 1$.

We will proceed as in Case II, except that here we analyze the expected change in POT^2 over two steps instead of one. The most desirable scenario is when station j crashes with station 2 in the first step while no station with $b_i > N$ attempts, and station 1 is the only station to transmit at the second step. Call this event E. Then

$$Ex[\Delta(POT^2)] = Ex[\Delta(POT^2)|E]Pr[E] + Ex[\Delta(POT^2)|\neg E]Pr[\neg E].$$

The decrease in potential when E happens is at least

$$(b_1+1)^{\alpha+1/2} - M - N^{\alpha+\frac{3}{2}} \geq \frac{1}{2}(b_1+1)^{\alpha+1/2}$$

provided $B \geq \max(M, N^2)$. The probability of E is at least

$$\frac{1}{2}p_j \frac{1}{(b_1+1)^{\alpha}} \prod_{i=2}^{N}(1-p_i')$$

where p_i' is the value of p_i after the first step as prescribed above. Reasoning as in Case II, we can then conclude that the first term is at most $-\frac{1}{4}(1-\lambda)(b_1+1)^{\frac{1}{2}}\frac{1}{M}2^{-N}POT$.

We have the same estimates for the second term as in Case II and this gives the desired conclusion provided that

$$B \geq \frac{c_B N^2 M^2 2^{2N}}{(1-\lambda)^2}$$

for some constant c_B independent of N and λ.

Case IV: We assume $p_1 = 1$; $\forall i > 1$ $p_i \leq \frac{1}{M}$.

In the last two cases, we can simplify the analysis for $Ex[\Delta(POT^2)]$ by finding bounds on $Ex[\Delta POT]$ and $Ex[(\Delta POT)^2]$ for some tree of descendent states. In particular, a simple calculation reveals that

$$Ex[\Delta(POT^2)] = 2POT \cdot Ex[\Delta POT] + Ex[(\Delta POT)^2]$$

for any set of descendent states and hence, we can prove that $Ex[\Delta(POT^2)] \leq -\delta POT$ by showing that $Ex[\Delta POT] \leq -\delta$ and $Ex[(\Delta POT)^2] \leq \delta POT$. We start by bounding $Ex[\Delta POT]$. In this case, we need only consider one step of the system. Proceeding as in Cases II and III, we find that $Q_i^+ = \lambda_i$,

$$Q_1^- \geq (1 - \frac{1}{M})^{N-1} \geq 1 - \frac{N}{M} \ ,$$

$$B_i^+ \leq \begin{cases} O(\frac{N}{M}) & \text{if } i = 1 \\ O(M^{-1/(2\alpha)}) & \text{if } i > 1 \end{cases}$$

and $B_i^- = 0$. Hence

$$Ex[\Delta POT] \quad \leq \quad \lambda - 1 + \frac{N}{M} + O(NM^{-1/(2\alpha)}) \leq -\delta$$

provided that $\lambda < 1$, $\delta < \frac{1-\lambda}{2}$ and

$$M \geq \frac{c_M N^{2\alpha}}{(1-\lambda)^{2\alpha}}$$

for some constant c_M independent of N and λ.

To finish the argument we estimate $Ex[(\Delta POT)^2]$ as follows:

$$\begin{aligned} Ex[(\Delta POT)^2] &\leq 4Ex[(\Delta Q^+)^2] + 4Ex[(\Delta Q^-)^2] + 4Ex[(\Delta B^+)^2] + 4Ex[(\Delta B^-)^2] \\ &\leq 4N^2 + 4 + O(\frac{N}{M} + NM^{1-1/\alpha}) \\ &\leq \delta POT \end{aligned}$$

for $POT \geq V$ where $V \geq \frac{c_V}{\delta} N^2 M$ and c_V is a constant that is independent of N and λ.

Case I: $\forall i$ $b_i \leq B$.

We will proceed as in Case IV. In particular, the bulk of the proof is devoted to showing that $Ex[\Delta POT] \leq -\delta$. Afterwards, we observe that $Ex[(\Delta POT)^2] \leq \delta POT$.

By making V to be large enough (the exact value will be determined later) and noting that if $POT \geq V$, we can assume that $q_N \geq \frac{V}{N} - (B+1)^{\alpha+1/2}$ and $b_N \leq B$ without loss of generality. In other words, the Nth queue is very large, and accounts for a good proportion of the overall potential.

The key to the proof is to show that with some not-too-small probability, the Nth station effectively dominates the channel for a very long time, thereby substantially reducing its massive queue and dramatically lowering the overall potential function. In particular, we show that there is a not-too-small probability that S_N is always the sole next station to broadcast after any collision. This is the hard part of the argument.

Once this is done, we finish up by showing that there aren't too many collisions over time and that not too many packets arrive over time. Of course, we must be sure to check that things can't get too bad if the Nth station ever does lose control.

To prove that we have a small probability of the Nth station staying in control we will first study what happens to a system of $N-1$ stations when we assume that all transmissions fail. This is essentially the situation when the Nth station never loses control.

Lemma 3.4: *Consider an isolated system where a single station advances from level i to $i+1$ with probability $i^{-\alpha}$ and otherwise remains at level i. Suppose the initial level of the station is between S and B. Then with probability exceeding $1-O\left(2^{-c\sqrt{S}}\right)$, the station reaches level b (for any $b \geq S$) within time $6\alpha b^{\alpha+1}$, and the station moves from level b to $b+1$ after time $\frac{b^{\alpha+1}}{4e^{\alpha+1}}$ for any $b \geq 2B$.*

Proof: Without loss of generality, we assume that the station starts at level S for the first part of the proof, and that it starts at level B for the second part. To avoid duplication of effort in the proof, we will use R to denote either S or B.

We start by computing the probability $Pr[b,t]$ that the transition from b to $b+1$ is made at step t. This probability is precisely

$$Pr[b,t] = \sum_{\substack{t_R+\cdots+t_b=t-r \\ t_j \geq 0}} \left(\prod_{j=R}^{b} (1-j^{-\alpha})^{t_j} j^{-\alpha} \right)$$

where $r = b+1-R$ and t_j denotes the number of steps that started and ended with the station in level j for $R \leq j \leq b$. Using the identity $1-x \leq e^{-x}$ and simplifying, we find that

$$Pr[b,t] \leq \prod_{j=R}^{b} \frac{1}{j^{\alpha}} \sum_{\substack{t_R+\cdots+t_b=t-r \\ t_j \geq 0}} e^{-\sum_{j=R}^{b} t_j/j^{\alpha}}$$

Since $\sum_{j=R}^{b} t_j = t-r$ and $t_j \geq 0$ for $R \leq j \leq b$, it is clear that $\sum_{j=R}^{b} \frac{t_j}{j^{\alpha}} \geq \frac{t-r}{b^{\alpha}}$.

Hence

$$\begin{aligned} Pr[b,t] &\leq \frac{(R-1)!^{\alpha}}{b!^{\alpha}} e^{-\left(\frac{t-r}{b^{\alpha}}\right)} \left(\sum_{\substack{t_R+\cdots+t_b=t-r \\ t_j \geq 0}} 1 \right) \\ &\leq \frac{(R-1)!^{\alpha} e^{r/b^{\alpha}}}{b!^{\alpha} e^{t/b^{\alpha}}} \binom{t}{r} \leq \frac{R^{\alpha R-\alpha} e^{\alpha b} e^{r/b^{\alpha}} t^r e^r}{b^{\alpha b} e^{\alpha R-\alpha} e^{t/b^{\alpha}} r^r} \leq \frac{e^{(\alpha+1)r} e^{r/b^{\alpha}} t^r}{b^{\alpha r} e^{t/b^{\alpha}} r^r} \\ &= \left(\frac{e^{\alpha+1} t e^{1/b^{\alpha}}}{b^{\alpha} r e^{t/rb^{\alpha}}} \right)^r \end{aligned}$$

In order to bound the behavior of this function, it is most useful to let $\beta = \frac{t}{b^{\alpha} r}$. Then

$$Pr[b,t] \leq \left[\frac{e^{\alpha+1} e^{1/b^{\alpha}} \beta}{e^{\beta}} \right]^r .$$

For large or small constant values of β, the preceding expression is very small. In particular, for $\beta \leq \frac{1}{2e^{\alpha+1}}$,

$$Pr[b,t] \quad \leq \quad \left(\frac{e^{1/b^\alpha}}{2e^{1/2e^{\alpha+1}}}\right)^r \quad \leq \quad \frac{1}{2^r}$$

assuming $b^\alpha \geq 2e^{\alpha+1}$ which will always be true since $b \geq R$. There are at most $\frac{b^\alpha r}{2e^{\alpha+1}}$ values of $t \leq \frac{b^\alpha r}{2e^{\alpha+1}}$. Hence, the probability that we progress from b to $b+1$ before step $\frac{b^\alpha r}{2e^{\alpha+1}} \leq \frac{b^{\alpha+1}}{2e^{\alpha+1}}$ is at most

$$\frac{b^\alpha r}{2e^{\alpha+1}2^r} \quad \leq \quad \frac{b^{\alpha+1}2^R}{2e^{\alpha+1}2^b} \quad \leq \quad O\left(\frac{b^{\alpha+1}2^R}{2^b}\right)$$

and thus with $R = B, b \geq 2B$ we have established the second part of the lemma.

For $\beta \geq 6\alpha$, the bound is at most

$$\left[\frac{6\alpha e^{1/b^\alpha}e^{\alpha+1}}{e^{6\alpha}}\right]^r \leq \frac{1}{2^r}\ .$$

This is small for $r \geq \sqrt{S}$. For smaller r we need to observe that $\beta = \frac{t}{b^\alpha r} \geq \sqrt{S}$ for $t = 6\alpha b^{\alpha+1}$, $b \geq S$ and $r \leq \sqrt{S}$ and in this case we use the bound

$$Pr[b,t] \leq \left[\frac{e^{\alpha+1}e^{1/b^\alpha}\beta}{e^\beta}\right]^r \leq e^{-c\sqrt{S}}$$

for some constant c. Moreover the bound forms a geometric series for $\beta \geq 6\alpha$, and thus the probability that the transition from b to $b+1$ is made after step $6\alpha b^\alpha r$ is at most $O(2^{-c\sqrt{S}})$ for $r \leq \sqrt{S}$ and for $r \geq \sqrt{S}$ we have the bound

$$O\left(\frac{b^\alpha r\alpha}{2^r}\right) \leq O\left(\frac{b^{\alpha+1}2^R}{2^b}\right)\ .$$

Summing over b again gives a geometric series, and for $r \geq \sqrt{S}$ we get the total estimate

$$1 - O\left(2^{-c\sqrt{S}}\right),$$

and the first part of the lemma is also established. ♠

Lemma 3.4 can be immediately extended to hold for $N-1$ isolated stations simultaneously by simply adding the failure probabilities. In other words, the result holds for $N-1$ stations simultaneously with probability exceeding $1 - O(N2^{-c\sqrt{S}})$.

Having established how the rest of the system behaves if the Nth station remains in control, we next look at the chances that the Nth station does maintain control.

Lemma 3.5: *Suppose S_N collides with another station at time T and backs off to $b_N = 1$. Then the probability that S_N will transmit successfully before any other station attempts a transmission is at least $1 - 2^\alpha \sum_{i=1}^{N-1} p_i$.*

Proof: Let $W = \prod_{i=1}^{N-1}(1-p_i)$ be the probability that none of the first $N-1$ stations try to send on a given step. Then the probability that the Nth station continues to maintain control after the collision is at least

$$2^{-\alpha}W + (1-2^{-\alpha})2^{-\alpha}W^2 + (1-2^{-\alpha})^2 2^{-\alpha}W^3 + \cdots =$$

$$\frac{2^{-\alpha}W}{1-(1-2^{-\alpha})W} = \frac{W}{2^{\alpha}-(2^{\alpha}-1)W}.$$

Replacing W with $1-\epsilon$, we observe that the probability of maintaining control is at least

$$\frac{1-\epsilon}{1+\epsilon(2^{\alpha}-1)} \geq 1-2^{\alpha}\epsilon\ .$$

Hence the probability of not regaining control at the next transmission is at most

$$2^{\alpha}\epsilon = 2^{\alpha}(1-W) \leq 2^{\alpha}\sum_{i=1}^{N-1} p_i$$

since

$$W = \prod_{i=1}^{N-1}(1-p_i) \geq 1-\sum_{i=1}^{N-1} p_i,$$

and the lemma follows. ♠

We now use Lemmas 3.4 and 3.5 to prove that S_N has a not-too-small probability of remaining in control for a very long time. The basic idea is that S_N keeps successfully transmitting until a collision occurs, whereupon it regains control before anyone else attempts to transmit. We will consider two kinds of collisions. The first involves collisions with stations that have backoff counters of size S or larger, and the second involves a collision with stations that have backoff counters of size less than S. There is also the possibility of both kinds of collisions happening simultaneously, but we will rig things so that this does not happen. By this we mean that the good set of events in which S_N remains in control this will not happen.

Collisions of the first kind are nice because the behavior of stations with backoff counters of size S or larger are governed by Lemma 3.4. Collisions of the second kind are nice because there are not very many of them, provided that we never allow the first $N-1$ stations to transmit successfully. In what follows, we consider sequences of events for which the Nth station always maintains control by directly blocking transmissions for other stations. We will show that no matter what times are chosen for the attempted transmissions of stations with small backoff counters, there is a not-too-small probability that everything works as we hope.

To start things off, we consider the probability that S_N succeeds in the first or second step. For this to happen we need S_N to broadcast at the first step to block anybody else from succeeding. We also keep anyone else from broadcasting at the second step so that S_N can succeed and establish control. This sequence of events happens with probability at least

$$\Omega(B^{-\alpha}(B+1)^{-\alpha}(1-\lambda)2^{-N}).$$

Henceforth, we will consider only sequences that started in this fashion, and thus have $b_N = 0$ at step 3.

Next define σ_γ to be the set of times (excluding steps 1 and 2) that one of the $N-1$ first stations would have made an attempt to transmit with backoff counter less than S if all its previous transmissions would have failed. Our argument will allow any possible configurations of σ_γ, observe only that by definition that $|\sigma_\gamma| \leq (N-1)S$. We partition σ_γ into $k \leq (N-1)S$ maximal intervals $I_1, I_2 \ldots I_k$ of configuration steps, and we define T_i to be the step following I_i for $1 \leq i \leq k$. By definition, $T_i \notin \sigma_\gamma$ for $1 \leq i \leq k$. Lastly set $\sigma'_\gamma = \sigma_\gamma \cup \{T_i | 1 \leq i \leq k\}$.

At each step of σ'_γ we will require that S_N attempts a transmission, and that each station with backoff counter S or larger does not attempt a transmission. This will insure that stations with small backoff counters never succeed, and that S_N regains control after a collision with any such station. Provided that the Nth station otherwise retains control (i.e. that $b_N \le 1$ before each I_i), the probability that these forced moves actually take place is at least

$$(2NS)!^{-\alpha}(1 - S^{-\alpha})^{2N^2S} \ge \frac{1}{2}(2NS)!^{-\alpha}$$

provided that $S \ge (4N^2)^{\frac{1}{\alpha-1}}$. The $(2NS)!^{-\alpha}$ factor is a gross underestimate on the probability that S_N transmits at all the desired times (which could all be bunched together in one large interval), and the $(1 - S^{-\alpha})^{2N^2S}$ factor accounts for the probability that the stations with large backoff counters do not attempt to transmit at all the desired times.

The preceding analysis accounts for collisions with stations that have small backoff counters. To account for stations that have large backoff counters, we apply Lemmas 3.4 and 3.5. In particular, we let E_t denote the event "At time t, S_N collides with another station, backs off to $b_N = 1$ and does not regain control by being the next station to send". If S_N loses control E_t must happen for some t. We need only analyze what happens outside σ'_γ, so S_N will only compete with stations with large backoff counters. We have by Lemma 3.5

$$Pr[E_t] \le 2^\alpha \left(\sum_{b_i \ge S} (b_i(t) + 1)^{-\alpha} \right)^2 \le 2^\alpha (N-1) \sum_{b_i \ge S} (b_i(t) + 1)^{-2\alpha}.$$

Where the second inequality follows by Cauchy-Schwartz inequality.

We next need to sum $Pr[E_t]$ over $t \notin \sigma'_\gamma$. To bound this probability, we will assume that the conclusion of Lemma 3.4 holds at time t but not necessarily at any future time, so as to avoid conditioning of the probabilities. We also have to be careful to note that the value of $b_i(t)$ depends on then S_i first had a backoff counter of size S, but otherwise is governed by Lemma 3.4. Combining these observations gives

$$\sum_{t \notin \sigma'_\gamma} Pr[E_t] \le 2^\alpha (N-1) \sum_{t \notin \sigma'_\gamma} \sum_{b_i \ge S} (b_i(t) + 1)^{-2\alpha}$$

$$\le 2^\alpha (N-1) \sum_{i=1}^{N-1} \sum_{t \notin \sigma'_\gamma, b_i(t) \ge S} (b_i(t) + 1)^{-2\alpha}$$

$$\le 2^\alpha (N-1)^2 \sum_{t=1}^{\infty} \left(max \left(S, \left(\frac{t}{6\alpha} \right)^{\frac{1}{\alpha+1}} \right) \right)^{-2\alpha}$$

$$\le 2^\alpha (N-1)^2 \left(\frac{6\alpha S^{\alpha+1}}{S^{2\alpha}} + \sum_{t=6\alpha S^{\alpha+1}}^{\infty} \left(\frac{t}{6\alpha} \right)^{\frac{-2\alpha}{1+\alpha}} \right)$$

$$\le c N^2 S^{1-\alpha},$$

assuming that the conclusion of Lemma 3.4 holds.

Although we still have many details to check, we are essentially done with the hard part of the analysis. In what follows we consider descendent states with depth at most $U + 2NS$ where $U + 2NS \leq q_N$. In particular, we are interested in sequences of descendent states for which the following conditions hold:

1) every backoff counter is at most $O(U^{1/(\alpha+1)})$,
2) S_N successfully broadcasts for all but $O(NU^{1/(\alpha+1)} \log U + NS)$ steps,
3) the number of new messages arriving overall in the first T steps is at most $\lambda T + 2NS + U^{1/2} \log U$ for all $T \leq U$, and
4) S_N gains control in the first 2 steps and maintains control thereafter (i.e. that the conditions described in the previous discussion are satisfied).

We first note that if all of these conditions hold for U steps, then we will have experienced a tremendous decrease in the potential function. This is because at least $U - O(NU^{1/(\alpha+1)} \log U)$ messages are successfully transmitted, at most $\lambda U + U^{1/2} \log U + 2N^2 S$ arrive, and each backoff counter adds at most $O(U^{\frac{\alpha+1/2}{\alpha+1}})$ to the potential. Hence, the decrease in potential is at least

$$(1-\lambda)U - O(NU^{\frac{\alpha+1/2}{\alpha+1}} + U^{1/2} \log U + N^2 S) =$$

$$(1-\lambda)U - O(NU^{\frac{\alpha+1/2}{\alpha+1}})$$

which is large for large U.

We next note that the probability that all of these conditions hold for U steps is not-too-small. This follows naturally form the proceeding analysis and Lemmas 3.4 and 3.5. In particular, the probability of gaining control in the beginning is

$$\Omega(B^{-2\alpha}(1-\lambda)2^{-N}).$$

Given that S_N gains control by the method described at the beginning, and that $b_N \leq 1$ at steps before I_i $(1 \leq i \leq k)$, the probability of having things go as planned for steps in σ'_γ is $\Omega((2NS)!^{-\alpha})$. Given that things have gone well at the beginning and during the previous steps of σ'_γ, the probability of not violating Lemma 3.4 nor having S_N otherwise lose control is at least

$$1 - O(N2^{-c\sqrt{S}} + N^2 S^{1-\alpha}).$$

The first term comes from Lemma 3.4 and the second comes from Lemma 3.5 and the above calculation. Thus we have calculated the probability of 4) holding. The probability of violating the first condition is $O\left(2^{-cU^{\frac{1}{2(\alpha+1)}}}\right)$ by Lemma 3.4. If the first condition is satisfied and S_N remains in control we know that there are at most $O(NU^{1/(\alpha+1)} + NS)$ collisions in the U steps. The time for the S_N to try to send again after each collision is at most $a \log U$ with probability $1 - O\left((NU^{1/(\alpha+1)}(1-2^{-\alpha})^{a \log U}\right)$. This is very small for $a = 2^{\alpha+1}$. Since the next attempted transmission will always be successful with probability $1 - O\left(\frac{N^2}{S^{\alpha-1}}\right)$, we can conclude that the Nth station successfully transmits on all but $O((NU^{1/(\alpha+1)} + NS) \log U)$ steps with probability $1 - O\left(\frac{N^2}{S^{\alpha-1}} + \frac{N}{U}\right)$. Hence with this probability condition 2 is satisfied. The last condition is easily verified to hold with probability $1 - O(1/U)$ by standard arguments.

Putting all the probabilities together, we find that all desired conditions hold with probability exceeding

$$\Omega(B^{-2\alpha}(1-\lambda)2^{-N}(2NS)!^{-\alpha})(1 - O(N2^{-c\sqrt{S}} + N^2 S^{1-\alpha} + \frac{N}{U}))$$

$$\geq \Omega(B^{-2\alpha}(1-\lambda)2^{-N}(2NS)!^{-\alpha})$$

for $U \geq c_U N$ and $S \geq c_S N^{\frac{2}{\alpha-1}}$.

Note that we have multiplied probabilities of success instead of adding probabilities of failure at two crucial points of the analysis. The first place we do this is at the beginning when we force S_N to gain control right away. Since later probabilities are conditioned upon this happening, multiplication of success probabilities for the first two steps and later steps is appropriate. The second place we multiply success probabilities is when we combine the probabilities that things work well during σ'_γ with the probability that things work well outside σ'_γ. Although these probabilities are not completely independent, the dependence is minimal and works in our favor. This is formally argued as follows.

Let $\rho_{i,t}$ and $\rho^*_{i,t}$ be random numbers drawn uniformly and independently from $[0,1]$ for $1 \leq i \leq N$ and $3 \leq t \leq U + 2NS$. The value of $\rho^*_{i,t}$ will be compared with λ_i to decide it the ith station gets a new packet at time t, and $\rho_{i,t}$ will be compared with $b_i(t)^{-\alpha}$ to decide if the ith station tries to transmit at the tth step. Note that the values of $b_i(t)$ depend on previous values of $\rho_{i,t}$ and $\rho^*_{i,t}$ for various i's and t's, but that the ρ- values are mutually independent.

For each S_i, examine the values of $\rho_{i,t}$ and $\rho^*_{i,t}$ for $3 \leq t \leq U + 2NS$ to determine the steps (if any) at which S_i would try to transmit with backoff counter less than S under the *assumption* (not the knowledge) that all attempts are blocked. Accumulating these values for $1 \leq i \leq N-1$ determine the time steps contained in σ'_γ. Note that the selection of steps that are in σ'_γ is independent of values of $\rho_{i,t}$ and $\rho^*_{i,t}$ for i such that $i = N$ or S_i has a backoff counter of size S or larger at step t under the assumption that all previous attempts have been blocked. Moreover, the remainder of this argument will not depend in any way on what steps were selected for σ'_γ.

We now analyze the probability that $\rho_{i,t}$ and $\rho^*_{i,t}$ are as we would hope for $t \in \sigma'_\gamma$ and i such that $i = N$ or S_i has a large backoff counter. In particular, we want S_N to try to broadcast for all $t \in \sigma'_\gamma$ and we do not want stations with a large backoff counter to try to transmit at any step $t \in \sigma'_\gamma$. As argued before, the ρ-values satisfy these demands with probability $\Omega((2NS)!^{-\alpha})$ provided only that $b_N \leq 1$ at the beginning of each I_i, $1 \leq i \leq k$. In addition, we can have at most N^2S new arrivals during σ'_γ, and S_N is thwarted from broadcasting for at most NS steps during these times.

We next consider the $\rho_{i,t}$ and $\rho^*_{i,t}$ values for $t \notin \sigma'_\gamma$ and i such that $i = N$ or S_i has a large backoff counter. To simplify the argument, we first consider the behavior of the system as in σ'_γ did not exist. In this scenario, we can apply Lemmas 3.4 and 3.5 and the analysis that followed to conclude that with probability $1 - O(N2^{-c\sqrt{S}} + N^2S^{1-\alpha} + \frac{N}{U})$, the values for $\rho_{i,t}$ and $\rho^*_{i,t}$ make the system perform exactly as desired. In other words, the values of the large backoff counters are regulated by Lemma 3.4, S_N never loses control, new packets do not arrive too fast, etc. Note that we add the probabilities of failure in this context, since the various modes of failure in the isolated system might be dependent. Also note that stations with small backoff counters are guaranteed not to attempt a transmission during these steps by the definition of σ'_γ.

Since the values of $\rho_{i,t}$ and $\rho^*_{i,t}$ for $t \in \sigma'_\gamma$ and $t \notin \sigma'_\gamma$ are independent we can conclude that all of the above constraints on the ρ- values are satisfied with probability

$$\Omega((2NS)!^{-\alpha})(1 - O(N2^{-c\sqrt{S}} + N^2S^{1-\alpha} + \frac{N}{U})).$$

Of course, we still must show what values of $\rho_{i,t}$ and $\rho^*_{i,t}$ that satisfy these constraints actually produce the desired sequence of events when we interleave steps in σ'_γ with steps not in σ'_γ in the correct order. Once this is accomplished, we are done since we will have shown that with probability exceeding

$$\Omega((2NS)!^{-\alpha})(1 - O(N2^{-c\sqrt{S}} + N^2S^{1-\alpha} + \frac{N}{U}))$$

the system behaves as claimed.

The proof that the real system behaves well if the $\rho_{i,t}$ and $\rho^*_{i,t}$ values satisfy the proceeding constraints proceeds by induction over t. The base case $t = 2$ was already established. We then consider what happens at some time $t \geq 3$, assuming that previous moves in the real system were essentially identical to moves of the corresponding steps of the isolated systems. If $t \in \sigma'_\gamma$, then we can be assured that $b_N \leq 1$ before the interval I_i that contains t, and thus that S_N broadcasts and that stations with big backoff counters do not broadcast. Hence, stations with small backoff counters are blocked and their behavior continues to agree with the assumptions that were used to define σ'_γ. Hence the definition of σ'_γ as describing when stations with small backoff counters attempt to broadcast remains valid. This is precisely what we want to have happen. For $t \notin \sigma'_\gamma$, then the system behaves exactly as it does in the scenario when we ignored σ'_γ because the activity in σ'_γ has no effect on any of the large backoff counters, and because stations with small backoff counters are guaranteed not to try anything by the definition of σ'_γ. The only possible difference is that b_N could be lowered from 1 to 0 by including some steps of σ'_γ. The only effect of making $b_N = 0$, however, is to start S_N broadcasting sooner. By the constraints, we know that S_N will be the next station to transmit anyway, so starting off sooner only increases the number of successful transmissions without otherwise changing that state of the system. Hence the behavior of the combined system is virtually identical to its behavior during σ'_γ and outside σ'_γ when considered in isolation, provided that the constraints on the ρ values are satisfied.

The formal justification that events go well with the claimed probability is now complete. All that remains is to bound the possible increase in the potential function should any of these conditions fail. Note that no matter what the reason for failure, all the conditions held in the previous step by assumption. Hence the most we could have added to the potential function because of the counters is $O(NU^{\frac{\alpha+1/2}{\alpha+1}})$. Similarly, the most we could add (net) because of the queues is $O(U^{1/2}\log U + NU^{1/(\alpha+1)}\log U)$. Hence, the worst increase we could suffer is $O(NU^{\frac{\alpha+1/2}{\alpha+1}})$.

Putting everything together, we find that there is a tree of descendent states with depth at most U for which the expected decrease in potential is at least

$$\Omega(B^{-2\alpha}(1-\lambda)2^{-N}(2NS)!^{-\alpha})\left((1-\lambda)U - O(NU^{\frac{\alpha+1/2}{\alpha+1}})\right) - O\left(NU^{\frac{\alpha+1/2}{\alpha+1}}\right)$$

$$\geq \Omega(B^{-2\alpha}(1-\lambda)^2U2^{-N}(2NS)!^{-\alpha}) - O\left(NU^{\frac{\alpha+1/2}{\alpha+1}}\right).$$

By selecting

$$U = \left(\frac{c_U 2^N NB^{2\alpha}(2NS)!^{\alpha}}{(1-\lambda)^2}\right)^{2\alpha+2}$$

for some constant c_U independent of N and λ, we get $Ex[\Delta POT] \leq -\delta$.

To complete the proof of Theorem 3.1, we need only observe that $Ex[(\Delta POT)^2]$ does not cause any problems since all the ΔPOT are of order U, and if we choose $V = \Omega(U^3)$ we are done.

Remark: Let us just point out that there is no problem in determining our constants. The reason is that our conditions can be summarized as follows.

$M \geq f_1(\lambda, N)$
$B \geq f_2(\lambda, N, M)$
$S \geq f_3(\lambda, N)$
$U \geq f_4(\lambda, N, B, S)$
$V \geq f_5(U, B)$
Here f_i are the explicit functions given in the proof.

4. Lower Bounds on $Ex[L_{ave}]$.

The analysis presented in Section 3 reveals that $\mathrm{Ex}[L_{ave}]$ is at most $P\left((1-\lambda)^{-1}\right) 2^{Q(N)}$ where P and Q are polynomial functions. We do not know whether or not the dependence on N can be made polynomial. The main difficulty in proving a polynomial upper bound in N by extending our methods lies in analyzing the probability that a station will grab control of the channel and empty its queue.

We can prove nontrivial lower bounds on $Ex[L_{ave}]$, however. In particular, in this section we show that for a wide range of backoff functions the expected number of nonempty queues over time is linear in the number of stations. For many backoff functions, this fact will imply that $Ex[L_{ave}]$ grows superlinearly in N.

We will assume that all stations have probability $\frac{\lambda}{N}$ of getting a message at each timeslot. This is a reasonable assumption since if the arriving messages are very unevenly distributed among the stations, the system would in reality be a system with fewer than N stations. On the other hand, small deviations from this assumption can be handled.

Let us start by giving an outline of the ideas of this section. The basic tool will be to establish a connection between our finite model and the infinite model briefly discussed in Section 2.3. We will not prove that the models behave in the same way, but rather that the proofs of instability in the infinite model extend to give lower bounds in our finite model. Before we can make this precise, however, we need a more formal definition of the infinite model.

In the infinite model, there is a countably infinite number of stations and no station ever gets two messages. The total number of messages that arrive to the system at a given time t is assumed to be Poisson distributed with mean λ. Each station behaves exactly in the same way as in the finite model. Since there are never two messages in any station, it is convenient to talk of the system as if each message had its own backoff counter. We next review some results for the infinite model.

Assume for the moment that the system is continuously externally jammed (i.e., that no transmissions are successful), and that a message arrives at time 1. Define $h(x)$ to be the probability that an attempt is made to transmit this message at time x. For example $Ex[h(x)] = \Theta(x^{-\frac{\alpha}{(\alpha+1)}})$ if $f(b) = (b+1)^{-\alpha}$. Let

$$H(\lambda) = \sum_{t=1}^{\infty} \left(1 + \lambda \sum_{x=1}^{t} h(x)\right) e^{-\lambda \sum_{x=1}^{t} h(x)}.$$

We will say that a backoff function has property H_λ if there exist λ' such that $\lambda > \lambda'$ and $H(\lambda')$ is finite. Using this notation we have the following theorem due to Kelly [5].

Theorem:*(Kelly [5]) Consider the infinite model and any backoff function satisfying property H_λ. Then expected number of successful transmissions up to time T is bounded by a constant independent of T, when the arrival rate is λ.*

It is not too difficult to check that essentially any function which grows slower than any exponential function has property H_λ for any $\lambda > 0$. In particular, this is true for any polynomial backoff function. The previous theorem does not apply to $f(b) = 2^{-b}$ for $\lambda < \ln 2$, however. In this case we need to rely on the following almost equally strong theorem.

Theorem: *(Aldous [1]) consider the infinite model and $f(b) = 2^{-b}$, $\lambda \leq \ln 2$. Then for any $a > 1 - \frac{\lambda}{\ln 2}$ the expected number of successful transmissions up to time T is $o(T^a)$.*

It will help to provide a brief outline for the proofs of these theorems. Let the *mass* $m(t)$ of the system at time t be defined by $m(t) = \sum_i f(b_i)$. A fact which underlies most of the analysis is that the probability of a successful transmission is bounded above by $(\lambda + m(t))e^{1-m(t)}$. An easy calculation shows that this is true in both the finite and infinite models.

Using this fact we can now give the idea behind the proofs. Within constant expected time, due to many messages arriving within a short time interval, the mass will exceed K for some large given constant K. Once this happens, the probability of success is small, and the messages that arrive to the system are not successfully transmitted, which implies that the mass increases even more, and so on.

To get a connection between the infinite and finite models, we will forget any message that is not an *active* message (i.e., first in its queue).

When we disregard messages which are not active, the finite model behaves in a similar way to the infinite model as long as the number of stations with empty queues remains at least cN. The only differences are

1) after a successful transmission, an additional new active message might appear in the finite model, due to the fact that the station in question has a queue of length 2 or more, and
2) the distribution of the number of arriving active messages is not constant over time in the finite model (it depends on the number of stations with empty queues) and is not Poisson (it is a sum of binomials instead).

Although these differences are substantial, for the most part they tend to just make the behavior of the finite system worse than its infinite counterpart as long as $x_e(t) = \Omega(N)$, where $x_e(t)$ is the number of empty queues in the finite system at time t. In particular we can establish the following key property.

Lemma 4.1. *Suppose that $x_e(t) \geq cN$ in an N-station system with arrival rates $\lambda_i \geq \frac{\lambda}{N}$. Then there a constants $K_{\lambda,c}$ and $d_{\lambda,c}$ such that within expected time $K_{\lambda,c}$ there is a point in time t_0 such that $m(t_0) \geq d_{\lambda,c}$ such that for every t*

$$m(t_0 + t) \geq \begin{cases} d_{\lambda,c} \log t & \text{for } f(b) = 2^{-b} \\ d_{\lambda,c} t^{\frac{1}{\alpha+1}} & \text{for } f(b) = b^{-\alpha} \end{cases}$$

or $x_e(t_0 + s) \leq cN$ for some $s \in [1, t]$

Proof sketch. Let us first take care of the case $f(b) = 2^{-b}$. Lemma 4.1 describes the mechanism that Aldous [1] uses in his proof. We will not repeat the proof here but just describe how to take care of the differences. Aldous uses two key lemma one which states that there are not too many successful transmissions (lemma 3) and one which states that there are many new arrivals (lemma 4). In our situation the proof of his lemma 3 goes through virtually without change. To prove the equivalent of lemma 4 one needs to take care of the differences described above. Difference 1) only helps us since it provides extra arrivals. To take care of difference 2) observe that what is needed is an estimate that much fewer messages arrives than expected. But since by our assumption on the number of empty queues and the arrival rates the expected value is high and the probability of getting only a fraction which is $\frac{7}{9}$ of the expected value is exponentially small also in our case.

The case $f(b) = b^{-\alpha}$ is easier and can be taken care of in two ways, either by imitating Aldous proof or by extending Kelly's proof to show that the expected number of transmissions before $x_e(t) \leq cN$ is a constant. The differences in the two models are taken care of in a similar way. ♠

Define

$$X_e(T) = \frac{1}{T}\sum_{t=1}^{T} x_e(t)$$

to be the average number of empty queues over time. We use Lemma 4.1 to bound $Ex[X_e(T)]$.

Theorem 4.2: *Let $f(b) = b^{-\alpha}$ or $f(b) = 2^{-b}$. Then for any $c > 0$ $E(X_e(T)) \leq cN + o(N)$, for $T \geq d_{c,\lambda}N$.*

Proof: Let $c' = \frac{c}{2}$. We know by lemma 4.1 that if $x_e(t) \geq c'N$ then within constant expected time the system will reach a state with $m(t) \geq -10 \log c\lambda$ and remain this way until $x_e(t) \leq c'N$. Once the mass is this large successes happen with probability $\leq \frac{c\lambda}{2}$. Since a message arrives in an empty queue with probability at least $c\lambda$ the number of empty queues will constitute a biased random walk. Furthermore the probability of not going into a high mass situation within time i is bounded by 2^{-ki} for some $k > 0$. This implies that if $x_e(0) \leq c'N$ then for any $t > 0$ $Pr[x_e(t) \geq c'N + i] \leq 2^{-ki}$ again with $k > 0$. If on the other hand $x_e(0) > c'N$ the probability that $x_e(t)$ remains greater than $c'N$ for time dN for some large d is less than $2e^{-d'N}$. This follows since with probability $e^{-d'N}$ the system will enter the state prescribed by lemma 4.1 before time $\frac{d}{2}N$ and the probability that the biased random walks stays above $x_e(t) \geq c'N$ is $e^{-d'N}$. From then on the previous case applies. In either case we know that for any $t \geq dN$ the probability that $x_e(t) \geq cN$ is $\leq e^{-d'N}$ and this proves the theorem. ♠

Having established that on the average we have a linear number of nonempty queues, we now look at the length of these queues.

Lemma 4.3: *If the system is stable and $f(b) = 2^{-b}$ or $f(b) = b^{-\alpha}$, then, for sufficiently large N, a fraction c_λ of the time $x_e(t) \leq \frac{2N}{3}$ and $m(t) \leq R_\lambda$.*

Proof: To have a stable system with total arrival rate λ, the probability of success has to be $\geq \frac{\lambda}{2}$ at least a fraction $\frac{\lambda}{2}$ of the time. This implies that $m(t) \leq R$ at least $\frac{\lambda}{2}$ of the time where R is a constant depending on λ. Furthermore, whenever $x_e(t) \geq \frac{2N}{3}$, by lemma 4.1 within constant expected time $m(t)$ will exceed R and stay that way for time at least cN. Thus the fraction of the time for which $m(t) < R$ and $x_e(t) \geq \frac{2N}{3}$ is bounded by $\frac{c}{N}$ and the lemma follows. ♠

Lemma 4.4: *If $f(b) = (b+1)^{-\alpha}$, $\alpha > 1$, then a fraction c_λ of the time there are $\Omega(N^{\frac{\alpha+1}{\alpha}})$ messages in the queues.*

Proof: We know by Theorem 3.1 that the system is stable and thus every state S of the system has a probability $Pr(S)$ associated which is the relative frequency with which the system is in state S. We know that

$$\sum_{S, m(S) \leq R, x_e(S) \leq \frac{2N}{3}} Pr(S) \geq c_\lambda.$$

For any state in the above sum there are $\Omega(N)$ queues whose backoff counters are at least $d_\lambda N^{\frac{1}{\alpha}}$. Define station i to be *unusual* if $\frac{\lambda b_i^{\alpha+1}}{4N(\alpha+1)} > q_i$ and $b_i \geq d_\lambda N^{\frac{1}{\alpha}}$ where d_λ is a constant to be determined. Say that any point in time is *unusual* if at least $\sqrt{N}$ of the stations are unusual. Using Lemma 3.4 it follows that the fraction of unusual points in time is exponentially small. Thus the states S which are not unusual and have $m(S) \leq R$ and $x_e(S) \leq \frac{2N}{3}$ have total probability $\geq c'$ for $N \geq N_\lambda$. But any such state has $\Omega(N)$ stations each with $\Omega(N^{\frac{1}{\alpha}})$ long queues and the lemma follows. ♠

Lemma 4.4 immediately implies the following lower bound on $Ex[L_{ave}]$.

Theorem 4.5: *If $f(b) = (b+1)^{-\alpha}$, then $Ex[L_{ave}] \geq \Omega\left(N^{\frac{\alpha+1}{\alpha}}\right)$.*

For quadratic backoff, this means that $Ex[L_{ave}] \geq \Omega(N^{\frac{3}{2}})$.

5. Instability of Exponential Backoff.

In this section, we prove instability results for binary exponential backoff. We will prove two results; one which is exact and the other asymptotic. Let us start by stating the exact result.

Theorem 5.1: *Suppose binary exponential backoff is used and the arrival rate at every station is $\frac{\lambda}{N}$ where $\lambda > \lambda_0 + \frac{1}{4N-2}$ and $\lambda_0 \approx 0.567$ is the solution to $\lambda_0 = e^{-\lambda_0}$. Then the system is unstable.*

To prove the result, we use the potential function

$$POT = C\sum_{i=1}^{N} q_i + \sum_{i=1}^{N} 2^{b_i} - N.$$

The best choice for C will turn out to be $2N-1$.

For any state in the system, we will show that the potential function is expected to increase by at least a fixed amount (independent of the state) during the subsequent transition. This will enable us to prove Theorem 5.1.

The proof requires the use of the following simple lemma.

Lemma 5.2: *If $0 \leq \epsilon_i \leq 1$ for $1 \leq i \leq m$, then*

$$\prod_{i=1}^{m}(1+\epsilon_i) \quad \geq \quad 1 + \sum_{i=1}^{m} \epsilon_i$$

and

$$\prod_{i=1}^{m}(1+\epsilon_i) \quad \geq \quad 2\sum_{i=1}^{m} \epsilon_i \ .$$

Proof: The first inequality is obvious from expansion of the product and the nonnegativity of the ϵ_i's. The second inequality follows from the observation that

$$\prod_{i=1}^{m}(1+\epsilon_i) - \sum_{i=1}^{m} 2\epsilon_i \quad \geq \quad \prod_{i=1}^{m}(1-\epsilon_i) \quad \geq \quad 0$$

since $\epsilon_i \leq 1$ for $1 \leq i \leq m$. ♠

In the proof, we let M denote the number of stations with a nonzero backoff counter. Without loss of generality, we can assume that $b_1, \ldots, b_M \neq 0$ and

$b_{M+1}, \cdots, b_N = 0$ where $0 \le M \le N$. Note that if $b_i \neq 0$, then $q_i \neq 0$ since there must be some message that failed in its most recent attempt to transmit. In addition, the queues in all but one of the stations $M+1, \dots, N$ must be zero. This is because any station with $b_i = 0$ and $q_i \neq 0$ must have successfully transmitted during the last step. Hence we divide our analysis into two cases, depending on whether or not $q_{M+1} = 0$.

Case I: $b_1, \dots, b_M \neq 0$; $b_{M+1}, \cdots, b_N = 0$; $q_{M+1}, \dots, q_N = 0$; $0 \le M \le N$.

We start with some additional notation. As in Section 3 we let

$$p_i = \begin{cases} 2^{-b_i} & \text{for } 1 \le i \le M \\ \lambda_i & \text{for } M+1 \le i \le N \end{cases}$$

be the probability that the ith station attempts to transmit where $\lambda_i = \frac{\lambda}{N}$ is the probability that a new message arrives at the ith station. Let

$$T = \prod_{i=1}^{N} (1 - p_i)$$

be the probability that none of the N stations attempts to transmit a message. In addition, the probability that none of $\{1, \dots, i-1, i+1, \dots, N\}$ try to transmit is $\frac{T}{(1-p_i)}$.

We also define

$$\epsilon_i = \frac{p_i}{1 - p_i}, \quad R = \prod_{i=1}^{N} (1 + \epsilon_i) \quad \text{and} \quad S = \sum_{i=1}^{N} \epsilon_i \ .$$

Note that $1 + \epsilon_i$ is $\frac{1}{1-p_i}$, so $RT = 1$. Also note that $0 \le \epsilon_i \le 1$ for $1 \le i \le N$ since $b_i \ge 1$ for $1 \le i \le M$ and $\lambda_i = \frac{\lambda}{N} \le \frac{1}{2}$ for $N \ge 2$.

We let Q_i^+, Q_i^-, B_i^+, and B_i^- denote the same quantities as in Section 3 and since expectations sum, we have

$$Ex[\Delta POT] = \sum_{i=1}^{N} Q_i^+ - \sum_{i=1}^{N} Q_i^- + \sum_{i=1}^{N} B_i^+ - \sum_{i=1}^{N} B_i^- \ .$$

It is easily seen that $Q_i^+ = C\lambda_i$ for $1 \le i \le N$. Since the ith station transmits successfully with probability $T\epsilon_i$, we can also easily conclude that $Q_i^- = CT\epsilon_i$ for $1 \le i \le N$. Since the ith station crashes with probability $\left(1 - \frac{T}{1-p_i}\right) p_i$, the value of B_i^+ is

$$\left(1 - \frac{T}{1 - 2^{-b_i}}\right) 2^{-b_i} \cdot 2^{b_i} = 1 - T(1 + \epsilon_i)$$

for $1 \le i \le M$, and

$$\left(1 - \frac{T}{1 - \lambda_i}\right) \lambda_i = \lambda_i - T\epsilon_i$$

for $M + 1 \le i \le N$. Finally, we note that

$$B_i^- = \frac{T}{1 - 2^{-b_i}} 2^{-b_i} (2^{b_i} - 1) = T$$

for $1 \le i \le M$, and that $B_i^- = 0$ otherwise.

Summing these values over $1 \le i \le N$,we find that

$$Ex[\Delta POT] = C\lambda - CTS + M - MT - TS + \sum_{i=M+1}^{N} \lambda_i - MT$$

$$= C\lambda + M + \sum_{i=M+1}^{N} \lambda_i - T[(C+1)S + 2M] \ .$$

Hence $Ex[\Delta POT] \ge \delta$ if and only if

$$R(C\lambda - \delta + M + \sum_{i=M+1}^{N} \lambda_i) > (C+1)S + 2M \ .$$

Substituting $C = 2N - 1$ and $\lambda \ge \frac{1}{2} + \frac{1}{2C} + \frac{\delta}{C}$, we need only check that

$$R(N+M) \ge 2NS + 2M$$

in order to verify that $Ex[\Delta POT] \ge \delta$ for any $\delta > 0$. This inequality easily follows from Lemma 5.2 since if $S \ge 1$, we use the facts that $R \ge 2S$ and $R \ge 2$, and if $S \le 1$, we use the facts that $R \ge 1 + S$ and $N \ge M$.

Case II: $b_1, \ldots, b_M \neq 0$; $b_{M+1}, \ldots, b_N = 0$; $q_{M+1} \neq 0$; $q_{M+2}, \ldots, q_N = 0$; $0 \le M < N$.

In this case, we are guaranteed that station $M+1$ will attempt a transmission. The probability that it is successful is

$$W = \prod_{i=1}^{M} (1 - 2^{-b_i}) \prod_{i=M+2}^{N} (1 - \lambda_i) \ .$$

The analysis for the Q_i's and B_i's is similar to Case I. In particular, $Q_i^+ = C\lambda_i$,

$$Q_i^- = \begin{cases} CW & \text{for } i = M+1 \\ 0 & \text{otherwise} \end{cases}$$

$$B_i^+ = \begin{cases} 1 & \text{for } 1 \le i \le M \\ 1 - W & \text{for } i = M+1 \\ \lambda_i & \text{for } M+2 \le i \le N \end{cases}$$

and $B_i^- = 0$ for $1 \le i \le N$. Summing these values, we find that

$$Ex[\Delta POT] = C\lambda - CW + M + 1 - W + \sum_{i=M+2}^{N} \lambda_i \ .$$

Thus $Ex[\Delta POT] \ge \delta$ if and only if

$$C\lambda + M + 1 + \sum_{i=M+2}^{N} \lambda_i \quad \ge \quad (C+1)W + \delta \ .$$

This is just a calculation and we defer it to the appendix.

Having established that we have an expected increase in potential each step let us see how we can use that to establish Theorem 5.1. There are general conditions under which increase in potential implies instability (see for instance [14]). However to verify these conditions require a fair amount of additional work and we believe that a direct proof is more illuminating.

Let us first prove that the expected waiting time is infinite over time.

Lemma 5.3: *At time T the expected waiting time for a newly arrived message is δT for some positive constant δ.*

A message that arrives at S_i has expected waiting time at least $q_i + 2^{b_i}$ the reason being that expected time before the first message in the queue is sent is 2^{b_i} and then at most one message can be sent per time step. The probability of an arriving message arriving at S_i is at least $\frac{\lambda}{N}$ if we assume that the expected number of arrivals per time step is less than 1 (otherwise the theorem is trivial). Thus the expected waiting time is at least

$$\frac{\lambda}{N}\sum_{i=1}^{N}(q_i + 2^{b_i}) \geq \frac{1}{4N^2}POT.$$

Since the expected value of the potential is $\Omega(T)$ we are done. ♠

Thus we have established that the expected waiting time and hence the expected queue size gets arbitrarily large as time goes by. Let us proceed to prove that the recurrence time is infinite.

Suppose we start at state where all queues are empty. We want to prove that the expected time to return to this state is infinite. We know that in time T the expected potential δT. As an extension of this we first establish that for some constants c and d it is true that with probability at least c the potential is at least dT. The essential lemma towards establishing this is:

Lemma 5.4: *For any $b \geq \lceil \log T \rceil$ the probability that i'th backoff counter has reached b in time T is bounded by $2^{-\binom{b-\lceil \log T \rceil}{2}}$.*

Proof: There must have been $b - \lceil \log T \rceil$ increases in the backoff counter after it reached $\lceil \log T \rceil$. There are

$$\binom{T}{b-\lceil \log T \rceil} \leq \frac{T^{b-\lceil \log T \rceil}}{(b-\lceil \log T \rceil)!}$$

possible ways to choose the time slots where these increases could happen. For any fixed choice of these time slots the probability that the backoff counter would increase at these time slots is at most

$$\prod_{i=\lceil \log T \rceil}^{b-1} 2^{-i} \quad \leq \quad T^{-(b-\lceil \log T \rceil)} 2^{-\binom{b-\lceil \log T \rceil}{2}}.$$

Multiplying out the lemma follows. ♠

Next we prove:

Lemma 5.5: *With probability at least c the potential at time T is at least dT.*

Proof: Let S denote a state of the system and consider the following claim:

Claim: There is a constant D such that

$$\sum_{S, POT(S) \geq DT} Pr[S]POT(S) \leq \frac{\delta}{2} T.$$

Before establishing the claim let us see how the lemma follows. Since $E(POT(S)) \geq \delta T$ the claim implies that

$$\sum_{S, POT(S) \leq DT} Pr[S]POT(S) \geq \frac{\delta}{2} T.$$

But this clearly implies that $Pr[POT(S) \geq \frac{\delta T}{4}] \geq \frac{\delta}{4D}$.

Thus we only have to establish the claim. Suppose $D > 4N^2$. Since no queue can be longer than T, the contribution from the queues to the potential is bounded by $2N^2T$. Thus for the potential to exceed DT it is necessary that $2^{b_i} \geq \frac{DT}{2N}$ for some i. Using that in this case the contribution to the potential from the queue lengths is bounded by the contribution from the backoff counters we get the estimate

$$\sum_{i=1}^{N} \sum_{S, b_i(S) \geq b_j(S), j \neq i, POT(S) \geq DT} Pr[S]POT(S) \leq N \sum_{b=\lceil \log \frac{DT}{2N} \rceil}^{\infty} 2N 2^b Pr[b_1 = b] \leq$$

$$2N^2 \sum_{b=\lceil \log \frac{DT}{2N} \rceil}^{\infty} 2^{b - \binom{b - \lceil \log T \rceil}{2}} \leq 4TN^2 \sum_{i=\lceil \log \frac{D}{2N} \rceil}^{\infty} 2^{i - \binom{i}{2}} \leq \frac{\delta}{2} T$$

for $D > D_\delta$. ♠

Now we are ready for the final part of the proof of Theorem 5.1. From Lemma 5.5 it follows that with probability at least $\frac{c}{T}$ the potential reaches at least dT before it returns to 0. Observe that the expected time to return from potential P to potential 0 is at least $\frac{P}{2N}$. This follows from looking at the largest backoff counter or the longest queue. Using this we get

$$Ex(return\ time) \geq \frac{1}{2N} Ex(Maximum\ potential\ before\ return)$$

$$= \frac{1}{2N} \sum_{i=1}^{\infty} i Pr[Max\ pot = i] = \frac{1}{2N} \sum_{i=1}^{\infty} Pr[Max\ pot \geq i] \geq \sum_{i=1}^{\infty} \frac{cd}{2Ni}.$$

However this last sum diverges and we have proved Theorem 5.1. ♠

Using the results for Section 4 we strengthen Theorem 5.1 slightly in an asymptotic sense.

Theorem 5.6: *Suppose binary exponential backoff is used and the arrival rate at every station is $\frac{\lambda}{N}$ where $\lambda \geq c$ where $c > \frac{1}{2}$. Then the expected recurrence time is infinite for $N > N_c$.*

Proof: Since the proof of Theorem 5.6 is almost identical to that of Theorem 5.1 we will only point out the modifications needed.

We will be working with the same potential function and will again show that the potential is expected to increase. However in this case we are sometimes forced to consider more than one step of the system to obtain the desired increase. We use the same cases as in the proof of Theorem 5.1. Observe first that in Case I we only needed $\lambda \geq \frac{1}{2} + \frac{1}{2C} + \frac{\delta}{C}$ to obtain the expected increase. Since $C > N$ this bound is $< c$ for $N > N_c$ and $\delta < 1$.

To handle Case II will require some work. We get two subcases depending on the number of stations with empty queues. Lemma 5.7 takes care of the first case.

Lemma 5.7: *For $c > \frac{1}{2}$ there is a constant d_c, $0 < d_c < 1$ such that if $\leq d_c N$ queues are nonempty the expected increase in potential is $> \delta$.*

Proof: We know by the analysis in Case II of Theorem 5.1 that the expected decrease if M stations have empty queues is given by

$$g(M) = 2N\lambda + (M+1)\left(1 - \frac{\lambda}{N}\right) - 2N\left(1 - \frac{\lambda}{N}\right)^{N-M-1}.$$

To prove that $g(M)$ is positive in the claimed interval we proceed as before by establishing that $g(N) > 0$, $g(d_c N) > 0$ and $\frac{\delta^2}{\delta M^2} g(M) < 0$. Since the first and last condition was taken care of in the proof of Theorem 5.1 we only need to establish the second condition. It is easy to see that $g(d_c N) > 0$ for sufficiently large N and $\lambda = c$ iff

$$\tilde{g}(d_c, c) = 2c + d_c - 2e^{-c(1-d_c)} > 0$$

But, since $\tilde{g}(d_c, c)$ is a continuous function in c and d_c and $\tilde{g}(1, c) > 0$ the lemma follows. ♠

To take care of the case of many empty queues we have.

Lemma 5.8: *Let $\lambda \geq c > \frac{1}{2}$, then there is a constant K_c such that if $x_e(t) \geq d_c N$ then the expected change in potential over the next K_c steps is $\geq \delta$.*

Proof: By the previous analysis the expected change in potential in a step is bounded from below by $-hN$ for some constant $h > 0$. Let $K' = \min(\frac{K}{2}, \frac{K}{32h})$ where K is a constant such that if $x_e(t) \geq (1 - d_c)N$ then the probability that $m(t + t_0) \geq 5$ for $t_0 \in [K', K]$ is $\geq 1 - \frac{K'}{K}$. Such a constant exist for sufficiently large N by lemma 4.1 applied with $c = \frac{1-d_c}{2}$. Consider the system over the next K steps. Observe that if $m(t+t_0) > 5$ then the expected increase in potential is at least $(c - \frac{1}{4})2N$. This follows since the probability of a successful transmission is $\leq \frac{1}{4}$ and the the contribution from the backoff counters is expected to increase. Let Δ_i be the expected change in potential in potential at time $t + i$. Then

$$\sum_{i=1}^{K} \Delta_i = \sum_{i=1}^{K'} \Delta_i + \sum_{i=K'+1}^{K} \Delta_i \geq$$

$$-hNK' + (1 - \frac{K'}{K})(K - K')(c - \frac{1}{4})2N - hN\frac{K'}{K}(K - K') \geq$$

$$NK(\frac{1}{8} - \frac{2}{32}) \geq \frac{NK}{16}.$$

This concludes the proof of lemma 5.8.

Using Lemmas 5.7 and 5.8 we know that the expected increase of potential over T steps is $\geq \delta T$ for some constant δ. We go from large expected potential to infinite recurrence time as was done in Theorem 5.1 and this completes the proof of Theorem 5.6. ♠

6. Instability of Linear and Sublinear backoff.

In this section we study linear and sublinear backoff and the goal of the current section is to prove the following theorem.

Theorem 6.1: *If $f(b) = (1+b)^{-\alpha}, 0 < \alpha \le 1$ is used as a backoff function then for any $\lambda > 0$ and $N > N_\lambda$ the system is unstable.*

As before we derive our result by using a potential function. In this case we use

$$POT = N^{\frac{3}{2}} \sum_{i=1}^{N} q_i - \sum_{i=1}^{N} (b_i+1)^{\alpha+1}.$$

We will first analyze the expected change in one or two steps. We do this by establishing a series of lemmas, and we start by giving some facts which are needed in several places. Let s be the number of nonempty queues. Let P_{suc} denote the probability of success. Then

$$P_{suc} = \left(1 - \frac{\lambda}{N}\right)^{N-s} \sum_{i=1}^{s} (1+b_i)^{-\alpha} \prod_{j \le s, i \ne j} (1 - (b_j+1)^{-\alpha})$$

$$+ \left(1 - \frac{\lambda}{N}\right)^{N-s-1} \frac{(N-s)\lambda}{N} \prod_{j=1}^{s} (1 - (b_j+1)^{-\alpha})$$

As in previous sections we need the expected change in the two components of POT. Observe that B^+ corresponds to the increase in POT and hence the decrease of $\sum_{i=1}^{N} b_i^2$. By straightforward analysis we have

$$Q^+ = \lambda N^{\frac{3}{2}}$$

$$Q^- = P_{suc} N^{\frac{3}{2}}$$

$$B^+ = \left(1 - \frac{\lambda}{N}\right)^{N-s} \sum_{i=1}^{s} (1+b_i)^{-\alpha}((b_i+1)^{\alpha+1} - 1) \prod_{j \le s, i \ne j} (1 - (1+b_j)^{-\alpha})$$

$$B^- = \sum_{i=1}^{s} (1+b_i)^{-\alpha} \left(1 - (1 - \frac{\lambda}{N}\right)^{N-s} \prod_{j \le s, i \ne j} (1 - (1+b_j)^{-\alpha}))((b_i+2)^{\alpha+1} - (b_i+1)^{\alpha+1}) +$$

$$\frac{(N-s)\lambda}{N} \left(1 - (1 - \frac{\lambda}{N})^{N-s-1} \prod_{j \le s} (1 - (1+b_j)^{-\alpha})\right)$$

B^- will not play any significant role in the analysis and the reason for this is the following fact:

Fact 1: $B^- \le 4N$

This follows from

$$B^- = \sum_{i=1}^{s}(1+b_i)^{-\alpha}\left(1-(1-\frac{\lambda}{N}\right)^{N-s}\prod_{j\le s, i\neq j}(1-(1+b_j)^{-\alpha}))((b_i+2)^{\alpha+1}-(b_i+1)^{\alpha+1})+$$

$$\frac{(N-s)\lambda}{N}\left(1-(1-\frac{\lambda}{N})^{N-s-1}\prod_{j\le s}(1-(1+b_j)^{-\alpha})\right)$$

$$\le \sum_{i=1}^{s}(1+b_i)^{-\alpha}((b_i+2)^{\alpha+1}-(b_i+1)^{\alpha+1}) + \frac{(N-s)\lambda}{N} \le 4s+N-s\le 4N$$

using $\frac{(x+1)^{\alpha+1}-x^{\alpha+1}}{x^\alpha} \le \frac{(\alpha+1)(x+1)^\alpha}{x^\alpha} \le (\alpha+1)2^\alpha \le 4$. The first inequality follows from taking the maximal value of the derivative and the last follows from $\alpha \le 1$.

Let us next take care of the easy case of estimating the change in potential.

Lemma 6.2: *If $P_{suc} < \frac{\lambda}{2}$ then $\Delta POT \ge \frac{\lambda}{3}N^{\frac{3}{2}}$ for $N > N_\lambda$.*

Proof: We have

$$\Delta POT \ge Q^+ - Q^- - B^- \ge \lambda N^{\frac{3}{2}} - \frac{\lambda}{2}N^{\frac{3}{2}} - 4N \ge \frac{\lambda}{3}N^{\frac{3}{2}}$$

for $N > N_\lambda$. ♠

In the future let c_λ be an arbitrary constant whose values depends on λ. We will assume that the value of c_λ may change from line to line and thus $2c_\lambda \le c_\lambda$ is a valid inequality. We are now considering $P_{suc} \ge \frac{\lambda}{2}$ and observe that this implies that $\sum_{i=1}^{s}(b_i+1)^{-\alpha} + (N-s)\frac{\lambda}{N} \le K_\lambda$, where K_λ is a constant close to $-\log\lambda$. Next we have

Lemma 6.3: *If $b_i > 0$ for $1 \le i \le s$ and $P_{suc} \ge \frac{\lambda}{2}$ then $\Delta POT \ge c_\lambda s^2 - N^{\frac{3}{2}} - 4N$.*

Proof: The main contribution to the increase of the potential this time will come from B^+.

$$B^+ = (1-\frac{\lambda}{N})^{N-s}\sum_{i=1}^{s}\frac{b_i^{1+\alpha}}{(1+b_i)^\alpha}\prod_{j\le s, i\neq j}(1-(1+b_j)^{-\alpha})$$

$$\ge c_\lambda \sum_{i=1}^{s}\frac{b_i^{2\alpha}}{(1+b_i)^\alpha} \ge c_\lambda s^2$$

.

The first inequality comes from $\sum_{i=1}^{s}(b_i+1)^{-\alpha} + (N-s)\frac{\lambda}{N} \le K_\lambda$ and the second inequality follows from Hölders inequality since

$$\frac{s}{2} \le \sum_{i=1}^{s}\frac{b_i^\alpha}{(b_i+1)^\alpha} \le \sum_{i=1}^{s}\frac{b_i^\alpha}{(b_i+1)^{\frac{\alpha}{2}}}\frac{1}{(b_i+1)^{\frac{\alpha}{2}}} \le \left(\sum_{i=1}^{s}\frac{b_i^{2\alpha}}{(1+b_i)^\alpha}\right)^{\frac{1}{2}}\left(\sum_{i=1}^{s}\frac{1}{(1+b_i)^\alpha}\right)^{\frac{1}{2}}$$

and again using that the last sum is bounded. The lemma now follows since $Q^- \le N^{\frac{3}{2}}$ and $B^- \le 4N$. ♠

Finally we must take care of the case when $b_1 = 0$ and $P_{suc} \ge \frac{\lambda}{2}$.

Lemma 6.4: *If $b_1 = 0$ and $b_i > 0$ for $2 \leq i \leq s$ and $P_{suc} \geq \frac{\lambda}{2}$ then $\Delta POT_{two\ steps} \geq c_\lambda s^2 - 2N^{\frac{3}{2}} - 8N$.*

Proof: The increase will come from B^+ in the case when a collision appears at step 1. The probability of collision at step 1 is

$$1 - \left(1 - \frac{\lambda}{N}\right)^{N-s} \prod_{i=2}^{s} (1 - (1+b_i)^{-\alpha})$$

$$\geq 1 - \prod_{i=2}^{s} (1 - (1+b_i)^{-\alpha}) \geq c_\lambda \sum_{i=1}^{s} (1+b_i)^{-\alpha}$$

since $\sum_{i=2}^{s} (1+b_i)^{-\alpha} \leq K_\lambda$.

If we have a collision at step 1 then by the proof of Lemma 6.3 at step 2 $B^+ \geq \sum_{i=1}^{s} c_\lambda \frac{b_i^2}{b_i+1}$. The value of b_i might have increased but since $\frac{b_i^2}{b_i+1}$ is increasing in b_i this would only make the inequality stronger. Thus the total expected increase in B over the two steps is at least

$$\left(c_\lambda \sum_{i=1}^{s} (1+b_i)^{-\alpha}\right) \left(c_\lambda \sum_{i=1}^{s} \frac{b_i^{2\alpha}}{(1+b_i)^{\alpha}}\right) \geq \frac{c_\lambda s^2}{4}.$$

Here we used the calculation done in the proof of Lemma 6.3. By the same estimates for Q^- and B^- as in the proof of Lemma 6.3 the lemma follows. ♠

Finally we will combine these results and with the aid of the results of section 5 obtain the instability of inverse backoff.

Lemma 6.5: *Let the system be at any state at time t. Then $E\left(POT(t + \frac{N}{10}) - POT(t)\right) \geq c_\lambda N^{\frac{5}{2}}$. for sufficiently large N.*

Proof: Observe that be the previous lemmas whenever $s > cN^{\frac{3}{4}}$ the expected increase per timestep in the potential is $\Omega(N^{\frac{3}{2}})$. To prove the lemma we need only establish that with high probability $s \geq cN^{\frac{3}{4}}$ during most of the interval. We have two cases. Remember that $s = N - x_e(t)$.

Case 1 $x_e(t) \leq \frac{4N}{5}$. Since at most one queue can become empty at each timeslot the number of nonempty queues remains large during the entire interval.

Case 2 $x_e(t) \geq \frac{4N}{5}$. By Lemma 4.1 it follows that for any $r \geq N^{\frac{4}{5}}$, $Pr[x_e(t+r) \geq N - \frac{\lambda r}{2}] \leq O(N^{-\frac{4}{5}})$. The reason being that to have many empty queues either there has been many successful transmissions or not too many messages have arrived. The probability of the first event is small by Lemma 4.1 and the second probability is easily seen to be exponentially small. Using this be get:

$$E\left(POT(t + \frac{N}{10}) - POT(t)\right) \geq \sum_{r=1}^{\frac{N}{10}} \Delta POT(t+r) \geq$$

$$N^{\frac{4}{5}} \times (-N^{\frac{3}{2}}) + \sum_{r=N^{\frac{4}{5}}}^{\frac{N}{10}} \Delta POT(t+r) \geq -N^{\frac{23}{10}} + c_\lambda N^{\frac{5}{2}} - c_\lambda N^{-\frac{4}{5}} N^{\frac{5}{2}} \geq c_\lambda N^{\frac{5}{2}}$$

and the lemma follows. ♠

Having established that the potential is expected to increase we now prove the theorem 6.1. Observe first that to prove that the expected queue size gets unbounded over time is trivial since the total queue size is always at least $\frac{POT}{N^{\frac{3}{2}}}$. We establish infinite expected recurrence time in the same way as in section 5. To make the same argument go through we only have to establish the lemma below.

Lemma 6.6: *Assume that $N > N_\gamma$ then there are constants c and d such that for each T, the probability that the potential at time T is at least dT is bounded from below by c.*

Proof: The expected potential at time T is δT. On the other hand

$$POT \leq N^{\frac{3}{2}} \sum q_i \leq N^{\frac{5}{2}} T.$$

Thus $Pr[POT > \frac{\delta T}{2}] \geq \frac{\delta T}{2N^{\frac{5}{2}}}$. ♠

Now the same argument as in Section 5 completes the proof of theorem 6.1. ♠

Remark Observe that in the proof of instability in the case of $f(i) = 2^{-i}$ we added a function depending on the backoff counters to the potential while in the case of $f(i) = (i+1)^{-1}$ we subtracted a function. This reflects a basic fact, namely that in the exponential case we back off too far while in the inverse case we back off too little.

7. Experimental Results.

We have simulated several backoff protocols for several different values of N and λ, and with different initial conditions. The experiments were done rather for exploratory reasons rather than a scientific investigation (with careful design and analysis of the experiments use of strong pseudorandom number generators etc.). We present here some of the results when the system starts with all empty queues. The values presented are for L_{ave} as computed over 10 million steps of the system. Each station was assumed to have arrival rate λ/N where λ varied between 0.1 and 0.8 and N varied between 2 and 300. We emphasize that we have only done one experiment for each set of parameter values and we have not done any careful analysis of the data. Thus we leave it to the reader to interpret the data in any way. We also encourage the reader to design a careful experiment to evaluate the various protocols in practice. We think this would be of great interest.

Table I

Observed values for L_{ave} after 10 million iterations of linear backoff.

N, λ	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8
2	0.02	0.12	0.44	1.4	5.5	24	97	420
5	0.03	0.24	1.1	9.0	$1.3 \cdot 10^5$	$6.4 \cdot 10^5$	$1.1 \cdot 10^6$	$1.6 \cdot 10^6$
10	0.04	0.29	1.9	$2.8 \cdot 10^5$	$7.8 \cdot 10^5$	$1.3 \cdot 10^6$	$1.8 \cdot 10^6$	$2.3 \cdot 10^6$
30	0.04	0.35	$4.4 \cdot 10^5$	$9.3 \cdot 10^5$	$1.4 \cdot 10^6$	$1.9 \cdot 10^6$	$2.4 \cdot 10^6$	$2.9 \cdot 10^6$
100	0.04	$3.7 \cdot 10^5$	$9.2 \cdot 10^5$	$1.4 \cdot 10^6$	$1.9 \cdot 10^6$	$2.4 \cdot 10^6$	$2.9 \cdot 10^6$	$3.4 \cdot 10^6$
300	0.04	$6.8 \cdot 10^5$	$1.2 \cdot 10^6$	$1.7 \cdot 10^6$	$2.2 \cdot 10^6$	$2.7 \cdot 10^6$	$3.2 \cdot 10^6$	$3.7 \cdot 10^6$

Table II

Observed values for L_{ave} after 10 million iterations of quadratic backoff.

N, λ	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8
2	0.04	0.31	1.4	6.5	26	79	230	810
5	0.06	0.51	3.1	26	160	610	1800	5800
10	0.07	0.55	3.6	51	840	19000	$1.3 \cdot 10^5$	$3.8 \cdot 10^5$
30	0.07	0.55	3.6	470	$3.8 \cdot 10^5$	$8.7 \cdot 10^5$	$1.4 \cdot 10^6$	$1.9 \cdot 10^6$
100	0.07	0.52	3.5	$3.4 \cdot 10^5$	$8.4 \cdot 10^5$	$1.3 \cdot 10^6$	$1.8 \cdot 10^6$	$2.3 \cdot 10^6$
300	0.07	0.53	3.5	$7.0 \cdot 10^5$	$1.2 \cdot 10^6$	$1.7 \cdot 10^6$	$2.2 \cdot 10^6$	$2.7 \cdot 10^6$

Table III

Observed values for L_{ave} after 10 million iterations of exponential backoff.

N, λ	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8
2	0.15	0.62	3.3	460	$2.3 \cdot 10^5$	$6.4 \cdot 10^5$	$1.1 \cdot 10^6$	$1.6 \cdot 10^6$
5	0.17	0.99	180	98000	$5.1 \cdot 10^5$	$7.1 \cdot 10^5$	$1.2 \cdot 10^6$	$1.7 \cdot 10^6$
10	0.17	0.99	96	$1.4 \cdot 10^5$	$3.5 \cdot 10^5$	$7.6 \cdot 10^5$	$1.2 \cdot 10^6$	$1.7 \cdot 10^6$
30	0.17	1.0	1200	$1.4 \cdot 10^5$	$4.9 \cdot 10^5$	$8.3 \cdot 10^5$	$1.3 \cdot 10^6$	$1.8 \cdot 10^6$
100	0.17	.95	610	$1.8 \cdot 10^5$	$5.1 \cdot 10^5$	$9.3 \cdot 10^5$	$1.4 \cdot 10^6$	$1.9 \cdot 10^6$
300	0.17	0.80	120	$1.9 \cdot 10^5$	$5.6 \cdot 10^5$	$1.0 \cdot 10^6$	$1.5 \cdot 10^6$	$1.9 \cdot 10^6$

8. Remarks.

A natural question is for what backoff functions can we prove Theorem 3.1. We believe that it can be extended to any natural backoff function $f(x)$ for which

$$\int_{x=1}^{\infty} f(x)dx < \infty$$

and

$$\frac{d}{dx}[f(x)^{-1}] << f(x)^{-1} \ .$$

Not that by integrating the second condition implies

$$f(x)^{-1} << \int_{b=1}^{x} f(b)^{-1} \ .$$

The first condition is necessary for proving that some station can dominate the channel for a long interval without ever losing control. Note that this condition is not satisfied for linear backoff, but is satisfied for more rapid backoff protocols. The second condition is necessary to argue that the expected benefit from decreasing a large backoff counter outweighs the expected damage from backing off the counter further. It is also necessary to insure that the other counters aren't backing off too far when one station is dominating the channel. Note that this condition is not satisfied for exponential backoff, but is satisfied for protocols that backoff less swiftly. A candidate for a potential function to use in such a stability result would be to pick a function $g(x)$ which grows faster than $f(x)^{-1}$ but slower than

$$\sum_{b=1}^{x} f(b)^{-1}$$

and then use a potential which is roughly

$$\sum_{i=1}^{n} q_i + \sum_{i=1}^{n} g(b_i) \ .$$

Based on our analysis, it would appear that the most popular and well studied protocols are precisely the wrong protocols. The good protocols, it would seem are the protocols in between.

Although we have made substantial progress in analyzing the performance of backoff protocols for communication in multiple access channels, we also leave several open questions. Most importantly, it would be nice to determine the behavior of $Ex[L_{ave}]$ as a function of N and $1-\lambda$ for polynomial backoff protocols. In particular, it would appear that the upper bounds are most in need of improvement. Once this is done, it might then be possible to decide which polynomial backoff protocol is best (i.e., which minimizes $Ex[L_{ave}]$ for a particular λ and N).

It would also be nice to completely determine the range of stability for exponential backoff.

9. Acknowledgements.

We are deeply indebted to Albert Greenberg for putting us straight with respect to the various definitions of stability. His comments saved us from potentially serious errors in the argument. We also want to thank Richard Koch for many discussions and suggestions and Richard Ladner, Robert Maier, Ron Rivest, Wojciech Szpankowski and John Tsitsiklis for their helpful comments and references. We are also grateful to one of the referees for pointing us to the book [8].

10. References.

1. D. J. Aldous, "Ultimate instability of exponential backoff protocol for acknowledgement based transmission control of random access communication channels," *IEEE Trans. on Information Theory,* Vol. IT-33, Number 2, March 1987, pp 219-223.
2. J. Goodman, A. Greenberg, N. Madras and P. March, "On the stability of the Ethernet," *Proc. 17th ACM Symp. on Theory of Computing,* 1985, pp. 379–387.
3. B. Hajek and T. van Loon, "Decentralized dynamic control of a multiaccess broadcast channel," *IEEE Trans. on Automatic Control,* Vol. AC-27, No. 3, June 1982, pp. 559–569.
4. *IEEE Trans. on Inform. Theory*, Vol. IT-31, March 1985.
5. F.P. Kelley, "Stochastic models of computer communications systems," *Journal of the Royal Statistical Society (B)* Vol. 47. No. 1, 1985.
6. R. Ladner, personal communication, 1986.
7. R. Metcalf and D. Boggs, "Ethernet: distributed packet switching for local computer networks," *CACM,* Vol. 19, No. 7, July 1976, pp. 395–404.
8. S.P. Meyn and R.L. Tweedie "Markov Chains and Stochastic Stability", Springer-Verlag, London 1993.
9. R. Rivest, "Network control by Bayesian broadcast," *IEEE Transactions on Information Theory*, Vol. IT-33, No. 3, May 1987, pp 323-328.
10. W. Rosenkrantz, "Some theorems on the instability of the exponential backoff protocol," in *Performance '84* (E. Gelenke, ed.), Elsevier, Amsterdam, 1984.
11. J. Shoch and J. Hupp, "Measured performance of an Ethernet local network," *CACM,* Vol. 23, No. 12, Dec. 1980, pp. 711–721.
12. Stidham, "The last word on the $L = \lambda W$," *Operations Research*, Vol. 22, 1974, 417-421.
13. W. Szpankowski, "Stability Conditions for multidimensional queuing systems with applications," *Operations Research*, Vol. 36, No. 6, 1988, 944-957.

14. W. Szpankowski and V. Rego, "Some theorems on instability with applications to multiaccess protocols", *Operations Research*, Vol 36, No. 6, November-December 1988, pp 958-966.

15. A. Tanenbaum, "Network protocols," *Computing Surveys,* Vol. 13, No. 4, Dec. 1981, pp. 453–489.

16. J. Tsitsiklis "Analysis of a Multiaccess Control Scheme," *IEEE Transactions on Automatic Control*, Vol. AC-32, No. 11, November 1987, pp 1017-1020.

17. B. Tsybakov and V. Mikhailov, "Ergodicity of a slotted Aloha system," *Problems of Information Transmission*, Vol. 15, No. 4, April 1980 (translated version).

A. Appendix. We do the calculation omitted is Section 5.

We need to prove that

$$C\lambda + M + 1 + \sum_{i=M+2}^{N} \lambda_i \quad \geq \quad (C+1)W + \delta \ .$$

Where quantities are as defined in the beginning of Section 5. Now we must make use of the fact that $\lambda_i \geq \frac{\lambda}{N}$ for all i. Setting $C = 2N-1$, it then suffices to prove that

$$(2N-1)\lambda + M + 1 + \frac{(N-M-1)\lambda}{N}$$
$$\geq 2N\left(1-\frac{\lambda}{N}\right)^{N-M-1} + \delta \ .$$

For $\lambda > \lambda_0 + \frac{1}{4N-2}$ where $\lambda_0 = e^{-\lambda_0} \approx 0.567$, the preceding inequality holds for any sufficiently small constant $\delta > 0$. To prove this, it is sufficient to show that $g(M) > 0$ for $0 \leq M \leq N-1$ and $\lambda > \lambda_0 + \frac{1}{4N-2}$ where

$$\begin{aligned} g(M) &= (2N-1)\lambda + M + 1 + \frac{(N-M-1)\lambda}{N} \\ &\quad -2N\left(1-\frac{\lambda}{N}\right)^{N-M-1} \\ &= 2N\lambda + (M+1)\left(1-\frac{\lambda}{N}\right) - 2N\left(1-\frac{\lambda}{N}\right)^{N-M-1} \ . \end{aligned}$$

We can then choose δ to be the minimum of $g(M)$ over $0 \leq M \leq N-1$.

We can show that $g(M)$ is always positive by proving that $g(0) > 0$, $g(N-1) > 0$ and that $\frac{d^2}{dM^2} g(M) < 0$ for $0 \leq M \leq N-1$. The only difficult part is showing that $g(0) > 0$ so we save it for last.

We start by showing that $g(N-1) > 0$. This is easy since

$$g(N-1) = 2N\lambda + N - \lambda - 2N = \lambda(2N-1) - N$$

is positive provided that

$$\lambda > \frac{N}{2N-1} = \frac{1}{2} + \frac{1}{4N-2} \ .$$

We next show that $\frac{\delta^2}{\delta M^2} g(M) < 0$ for $0 \le M \le N-1$. This is also easy since

$$\frac{\delta}{\delta M} g(M) = \left(1 - \frac{\lambda}{N}\right) + 2N \ln\left(1 - \frac{\lambda}{N}\right)\left(1 - \frac{\lambda}{N}\right)^{N-M-1} .$$

and

$$\frac{\delta^2}{\delta M^2} g(M) = -2N\left(\ln\left(1 - \frac{\lambda}{N}\right)\right)^2\left(1 - \frac{\lambda}{N}\right)^{N-M-1} .$$

Lastly, we must prove that $g(0) > 0$. The argument here is a bit trickier because we must be careful when bounding the $(1 - \frac{\lambda}{N})^{N-1}$ term in the expression for

$$g(0) = 2N\lambda + 1 - \frac{\lambda}{N} - 2N\left(1 - \frac{\lambda}{N}\right)^{N-1} .$$

We start by observing that

$$(1 - \frac{\lambda}{N}) = e^{-\frac{\lambda}{N} - \frac{\lambda^2}{2N^2} - \frac{\lambda^3}{3N^2} - \cdots}$$

and thus that

$$\begin{aligned}
(1 - \frac{\lambda}{N})^{N-1} &= e^{\ \lambda + \frac{\lambda}{N} - \frac{\lambda^2}{2N} + \frac{\lambda^2}{2N^2} - \frac{\lambda^3}{3N^2} + \frac{\lambda^3}{3N^3} - \cdots} \\
&= e^{-\lambda + \frac{\lambda(2-\lambda)}{2N} + \frac{\lambda^2(3-2\lambda)}{6N^2} + \frac{\lambda^3(4-3\lambda)}{12N^3} + \cdots} \\
&\le e^{-\lambda + \frac{1}{2N} + \frac{1}{6N^2} + \frac{1}{12N^3} + \cdots} \\
&< e^{-\lambda + \frac{1}{2N} + \frac{1}{6N^2}\left(\frac{1}{1-1/N}\right)} \\
&\le e^{-\lambda + \frac{1}{2N-1}} \\
&= e^{-(\lambda - \frac{1}{4N-2}) + \frac{1}{4N-2}} \\
&\le \left(\lambda - \frac{1}{4N-2}\right) e^{\frac{1}{4N-2}} \\
&\le \left(\lambda - \frac{1}{4N-2}\right) \frac{1}{1 - \frac{1}{4N-2}} \\
&\le \lambda
\end{aligned}$$

since $\lambda - \dfrac{1}{4N-2} > \lambda_0$ and $e^{-\lambda'} \le \lambda'$ for all $\lambda' \ge \lambda_0$, and $e^x \le \dfrac{1}{1-x}$ for $0 \le x < 1$. Hence

$$g(0) \ge 2N\lambda + 1 - \frac{\lambda}{N} - 2N\lambda > 0 .$$