




**Stability and Performance of Contention
Resolution Protocols**

by

Hesham M. Al-Ammal

Thesis

Submitted to the University of Warwick

for the degree of

Doctor of Philosophy

Computer Science

August 2000



Contents

Acknowledgments	iv
Declarations	v
Abstract	vi
Chapter 1 Introduction and Related Results	1
1.1 The Multiple-Access Channel	2
1.2 The Protocols	4
1.3 Performance Measures	6
1.4 Related Results	9
1.4.1 The Infinite Model	9
1.4.2 The Finite Model	12
1.5 Which protocols are we interested in?	15
1.6 Summary of Results	18
Chapter 2 Stability Conditions for Markov Chains	21
2.1 Definitions and Notation	22
2.2 Definitions of Stability	24
2.3 Establishing Positive Recurrence and Strong Stability Using Drift	27
2.3.1 Multiple step drifts	28
2.3.2 Strong stability	29

Chapter 3	The Stability of c-ary Exponential Backoff	38
3.1	The Model and the Protocol	39
3.2	Choosing a Suitable Potential Function	40
3.3	The Stability of c -ary Exponential Backoff	44
Chapter 4	Instability Results for Exponential Backoff Protocols	63
4.1	Instability of c -ary Exponential Backoff	64
4.2	Instability of the Two-User Binary Exponential Backoff	72
4.2.1	Bounded Expected Jumps	75
4.2.2	Non-negative drift	76
Chapter 5	Age-based Protocols in the Finite Model	82
5.1	Strong Stability of an Age-based Protocol Similar to Binary Exponential Backoff	83
5.1.1	Conditions (C1) and (V2)	86
5.1.2	Negative drift	92
5.2	The Stability of Sublinear Age-based Protocols	102
Chapter 6	Remarks and Open Problems	108
6.1	Remarks and Open Problems on the Protocols	109
6.1.1	Exponential Backoff Protocols	110
6.1.2	Age-based Protocols	113
6.1.3	The Search for n	115
6.2	Other Techniques for Showing Stability	117

Acknowledgments

First and foremost, I would like to thank my thesis supervisor Dr. Leslie Ann Goldberg for all her help, advice, and encouragement throughout my studies at Warwick. I have truly learned a great deal from her and from the countless discussions at her office. Her love for research and science has been an inspiration to me, and her knowledge and wisdom guided me through the complex and hard paths toward completing this thesis.

I would like to thank my wife Wafa, for her constant encouragement, help, understanding, and devotion. Without her I couldn't have reached so far. I thank my parents for their unconditional support and for giving me an environment which made me love science.

Finally, my thanks goes to the administration of the University of Bahrain, and to my colleagues at the computer science department; who supported my post-graduate studies.

Declarations

This thesis is submitted to the University of Warwick in support of my application for admission to the degree of Doctor of Philosophy. The material in this thesis is my own work, except where it is acknowledged otherwise. No part of this work has been submitted in support of an application for another degree or qualification in this or any other institution of learning. A special case of the result presented in Section 3.3 appeared in the following refereed paper in which my own work was that of a full pro-rata contributor:

Al-Ammal, H. , Leslie Ann Goldberg, and Phil MacKenzie, *Binary Exponential Backoff is Stable for High Arrival Rates*, Proceedings of the International Symposium on Theoretical Aspects of Computing 2000, vol. 17, pp. 169-180, Springer Verlag. (The same article is to appear in the journal *Theory of Computing Systems* under the title: An Improved Stability Bound for Binary Exponential Backoff).

Hesham M. Al-Ammal,
August 2000

Abstract

The stability of contention resolution protocols for a multiple-access channel is investigated. We start by outlining the different definitions used for stability of such protocols. Then we present sufficient conditions for the positive recurrence of countable state space Markov chains using a potential function defined on their state space. By adding a new condition to the usual negative drift requirement, we show a stronger form of stability which ensures that the expected value of the potential is bounded in the steady state. These conditions are then used to prove stability for various contention resolution protocols in the time-slotted, Bernoulli arrivals, finite model. We study several acknowledgement-based protocols under this model because of their practicality and the minimal assumptions required by this class of protocols.

Let n be the number of users in the system, and suppose that each user has a mean arrival rate λ_i , and the total system arrival rate is $\lambda = \sum_{i=1}^n \lambda_i$. We start by examining the c -ary exponential backoff protocol which works as follows. For any constant $c > 1$, user i with a message that encountered b_i collisions so far, transmits with probability c^{-b_i} . This class includes binary exponential backoff (a variant of the Ethernet protocol), and we show that for symmetric arrivals and sufficiently large n , c -ary exponential backoff is positive recurrent for any $\lambda < 1/(\alpha n^{3/4+\epsilon})$ for some constant α , and any $\epsilon > 0$. Furthermore, we prove for $c \geq 2$ that the protocol is unstable (i.e. not positive recurrent) for any $\lambda > \lambda_0 + 1/(4n - 2)$, where λ_0 is the solution to the equation $\lambda_0 = e^{-\lambda_0}$. For the special case when $n = 2$ and $c = 2$, we present new conditions on the arrival rates which are sufficient for instability.

Next, the stability of age-based protocols (also acknowledgement-based) is studied. In age-based protocols, a message waiting at the head of the queue for t steps is transmitted with probability p_t . We show that the age-based protocol where $p_t = 4/(t+4)$ is strongly stable (in the sense that there exists a stationary distribution and the expected load of the system is finite) for symmetric arrivals and any arrival rate $\lambda = 1/(\alpha n^{3/4})$. We also show that any age-based protocol sending with probability $p_t = t^{-\epsilon}$ for some $0 < \epsilon < 1$ is positive recurrent for any $\lambda = O(n^{-1-1/\epsilon})$. Finally, we conclude with some remarks and open questions.

Chapter 1

Introduction and Related Results

Consider the problem of sharing a single resource by a distributed set of users. These users have no way of communicating or coordinating their attempts to use this resource, and when two or more users attempt to access this resource at the same time none succeed and they must try again at a later time. The users can, however, utilise the information obtained from the availability of the shared resource to schedule their future attempts. A solution to this problem may be obtained by creating a set of rules that allow the users to share the resource efficiently and with a minimum number of failed attempts. Intuitively, these rules are often stochastic in nature because of symmetry breaking reasons. This set of rules is called a *contention resolution protocol*, and this is the classical problem of sharing a multiple-access broadcast channel. Even though the shared resource is modelled by a broadcast communication channel (which is mainly due to historical reasons), contention resolution protocols have found many other applications ranging from the Internet to parallel optical computing.

In this thesis we will study *simple* contention resolution protocols which make a minimal set of assumptions about the kind of feedback received from

the shared resource and do not have knowledge of the total number of users. Members of this class of protocols are said to be *acknowledgement-based*, and a precise definition of this class will be presented later. We will mainly be interested in answering the question: “under which conditions does a protocol work well?” By “conditions” we generally mean a maximum allowable arrival rate for the requests made by the users. Throughout this thesis we will denote the mean arrival rate of a system by λ . A protocol works “well” when it is able to satisfy these requests in a reasonable (or at least not infinite) amount of time. Following past practice in the research literature, the shared resource is modelled by a multiple-access broadcast channel, and requests are modelled by incoming messages.

We start this introductory chapter by defining formally this multiple-access channel. We also present the properties and assumptions for two important models for this channel, namely the finite and the infinite models. Then we describe different classes of contention resolution protocols, classified according to the amount of feedback and knowledge necessary for each protocol to work correctly. We will also list some performance measures which were adopted by researchers for picking a “good” protocol. A review of past research in the area of contention resolution protocols is then presented, covering both models. We finish this introductory chapter by a summary of the results of the thesis.

1.1 The Multiple-Access Channel

In the physical world, a multiple-access channel is a broadcast channel such as a coaxial cable, fiber optic wire, radio channel, or satellite channel. This channel is used to provide a means of communication between a geographically distributed set of users with no other means of communication. The channel can also be a metaphor for any other resource shared between a distributed set of users. Because of the broadcast nature of the channel, when two or more users attempt to communicate, their transmissions are scrambled and they must at-

tempt to send their messages at a later time.

As a mathematical model, the *multiple-access channel* is a single noiseless channel shared among a set of distributed users. Time is divided into discrete steps, or slots. Users communicate by sending messages through the channel. Message sizes are fixed and a single message can be transmitted in a single time unit. Communication over the channel is not centrally controlled, and thus a *contention resolution protocol* is needed to synchronise transmissions. Users of the channel can detect at the end of a transmission whether it was successful or not. Any step where more than one user sends is called a *collision*. There are two basic models for the channel:

1. **The infinite model.** This model has appeared under several names in the literature, including the infinite population, Poisson [28, 25], and Queue-free model [20, 21]. In this model users do not maintain queues, and every arriving message enters as a new user. At the beginning of every time step, a number of new users enters the system. This number is drawn from a Poisson distribution with mean λ . All of the messages present in the system participate in the contention resolution process. Thus, each message decides independently of all others whether or not to send during the current step. This decision is controlled by the specific contention resolution protocol used in the system. In the event that a single message sends during a time step, it is considered successful and is removed from the system. If more than one message sends, a collision is detected and the messages must be sent again at a later time.
2. **The finite model.** Also called the Bernoulli [28, 25] or Queueing model [20, 21]. In this model there are exactly n users sharing the multiple-access channel. Each user maintains a queue of infinite capacity which holds the arrivals into the user. Arriving messages are added to the queues of the users, and at each step user i receives a single new arrival with probability λ_i . The total mean arrival rate is taken to be $\lambda = \sum_{i=1}^n \lambda_i$. During each time step a user decides whether or not to send the message at the head

of its queue independently of other users. If exactly one user sends during the current step, the message is successfully sent and is removed from the queue. Otherwise, all messages remain in the queues and must await retransmission at a later time.

There are advantages for studying a specific protocol under each of the two models described above. The infinite model enables us to better understand the contention process and isolates it from the queueing process. On the other hand, the finite model is more realistic when modelling real applications such as local area or satellite networks. Furthermore, the queueing process seems to have a stabilising effect which results in higher arrival rates for the finite model (as we will see later in the review of related results).

Many authors favour working with the infinite model because they consider it as a “worst case” situation for the finite version of the protocol. They suggest that any protocol which is proven to be stable for the infinite model can be transformed into a similar protocol that is stable for the finite model for some (possibly smaller) arrival rate. This transformation enables the finite model protocol to simulate the infinite one [43]. In this simulation, each message in a finite model system will act as a user and will participate in the contention resolution process (even those waiting in the queues). However, these arguments (which date back to Abramson [2]) are heuristic in nature and give no formal method for the simulation or calculating the new stable arrival rate.

1.2 The Protocols

The research into protocols for sharing a multiple-access channel has a rich history in the Electrical Engineering and Computer Science literature [39]. Most of this research focuses on simple protocols because they are easier to implement and understand. Many of the applications require the protocol to be as simple as possible because of space limitations in the communicating devices. Also, a simple protocol is faster to execute and creates less overhead per time

step.

The amount of channel sensing done by the users is an important aspect of any protocol. Thus, a natural way of classifying contention resolution protocols is by the kind of feedback required by the protocols to determine their transmission schedule. A user in a protocol may get *binary*¹ feedback in the form of an acknowledgement if it sends successfully, or a non-acknowledgement when sending its transmission coincides with others sending at the same step. In the *ternary* feedback model (also called 0-1-e feedback), the user receives 0 when all users are silent, 1 when exactly one user sends, or e (for error) when more than one user sends to the channel. A more powerful kind of sensing is *m*-ary feedback, where all users receive a non-negative integer *m*, where *m* is the number of users that sent a message during the previous step.

Contention resolution protocols can also be divided into the following classes according to the amount of information needed by the protocol to decide when will the user send. Because of the difference in power available to each class of protocols, this classification will help us while comparing previous results according to the assumptions made by each protocol. Note that a user is called *active* when it has at least one message waiting to be sent.

1. **Full-sensing protocols:** Every user (active or not) in this class of protocols listens to the channel in every step. All users get 0-1-e feedback (or *m*-ary feedback in more powerful models) after every step.
2. **Acknowledgement-based protocols:** In acknowledgement-based protocols [28], an active user has knowledge of the history of its own transmission attempts only. This class was introduced because in many applications it is impossible to listen to the channel all the time. Hence, in this class a user only needs to know whether its own transmissions were successful or not. This is usually done by the user receiving acknowledge-

¹It should be pointed out that the term has different meanings in the multiple-access protocols literature. Some authors use it mean that the users get a *success* or *failure* signal, where the station is notified of its attempts only. Others use the term to mean *conflict* and *no conflict*, where no conflict may mean success or idle.

ments from the channel when its transmission is successful. It is also assumed that the user in this class has no knowledge of the number of users in the system. There are two popular subclasses of acknowledgement-based protocols:

- (a) **Age-based Protocols:** An age-based protocol [34] associates a probability $p_0, p_1, p_2, \dots$ for sending a message by an active user with the age of this message. Every active user with a message waiting to be sent for i steps will attempt to send it independent of other users with probability p_i .
 - (b) **Backoff protocols:** Active users in backoff protocols maintain a counter of the number of collisions encountered by a message waiting to be sent. All users with collision (or backoff) counter equal to i send independently with probability p_i . There are many popular examples of backoff protocols, including the Ethernet [40] and ALOHA [1] protocols.
3. **Almost-acknowledgement-based protocols:** These are protocols that require in addition to the history of the transmissions of active users, a rough estimate on the total number of users in the system [46, 23].

1.3 Performance Measures

There are many ways of measuring how “good” a protocol is. The most basic property that should be present in a good protocol is *stability*. The intuitive idea behind stability is that the protocol should deliver the incoming messages, and that as time goes by, the messages present in the system should not increase in an unbounded manner. An obvious measure that a protocol designer might be interested in is the *average system load* L_{avg} , which is defined as the average number of messages waiting in the system. We can also measure W_{avg} which is defined as the average number of steps a message has to wait before it is successfully delivered to its destination. Another performance measure is the

return time to the empty state T_{ret} , which is defined as the time needed to return to the start state (the state where there are no messages in the system).

Our main tool for analysing these contention resolution protocols will be the Markov chain X_t which will model the behaviour of the system. Each state of this Markov chain is usually specified by the queue sizes of the users and any other relevant information used by the protocol to decide when to transmit messages. A reasonable assumption about all of the systems studied here is that all users start with an empty queue. Thus, we assume that the chain has a unique state $X_0 = (0, 0, \dots, 0)$ which is called the *empty state*. Basically, in this state all users have empty queues. For the chain X_t (and the protocol modelled by X_t), define the return time to the start state as

$$T_{ret} = \min\{t \geq 1 \mid X_t = X_0\}.$$

The *system load* at step t , denoted by $L(X_t)$, is defined as the total number of messages waiting to be sent (i.e. the sum of the queue sizes) during step t . The average load of the system is defined explicitly as

$$L_{avg} = \lim_{t \rightarrow \infty} \frac{1}{t} L(X_t).$$

Note that the random variable L_{avg} may not exist, but we are only interested in situations where we can prove that it does. Another way of characterising the system load is by examining it in the steady state of the system. If we can show that the system has a stationary distribution π , we will also be interested in showing that $E_\pi[L(X_t)]$ is finite. We use the notation $E_\pi[x]$ to denote the expected value of the random variable x in the distribution π .

In the literature, there are many definitions of stability, and we will precisely describe the definitions adopted for this thesis later in Chapter 2. However, we will define briefly here some popular notions of stability to help us in reviewing past literature. Often, when the arrivals to the users are chosen according to a fixed well defined distribution, the evolution of a protocol can be modelled by a Markov chain. Most of these definitions of stability and instability depend on this underlying Markov chain and its stability.

Next, we present some popular definitions of stability:

1. **Recurrence:** This is basically a property of the underlying Markov chain. An irreducible and aperiodic Markov chain is said to be recurrent if, starting from any state, the probability that it will return to the empty state is one.
2. **Positive recurrence:** A protocol is said to be positive recurrent (some times called ergodic) if, starting from the empty state, $E[T_{ret}] < \infty$. This is again a property of the underlying Markov chain which is assumed to be irreducible and aperiodic, and it implies the existence of a unique stationary distribution for the chain.
3. **Bounded expected load:** Since buffer size is a major concern for hardware designers, we may require that the load be bounded in some way. This certainly has been a major concern for researchers in adversarial queueing networks [6], where the authors define a system to be stable when the expected number of messages waiting to be delivered is finite ² (In particular, they require $\sup_t E[L(X_t)] < \infty$, and say nothing about the process having a stationary distribution or not).
4. **Strong stability:** A protocol is strongly stable if, starting from the empty state, both the expected return time to the empty state is finite, and the expected average load is finite. This strong form of stability was proposed by Håstad, Leighton, and Rogoff [27] in an attempt to unify the different notions of stability in a single definition which encompasses the others.

Next, we turn our attention to definitions of instability. A protocol is said to be *transient* if it is not recurrent (and hence the probability of returning to the empty state is less than 1). A protocol is said to be *unstable* if it is not positive recurrent (hence the expected return time to the empty state is infinite). Note that transience is a stronger form of instability than non-positivity [41].

²To a probabilist, this notion of stability is strange [45]. They usually think of stability as the existence of a unique stationary distribution for the chain. However, to a system designer, the main concern is having a bounded number of messages in the system to ensure that the buffers are not overloaded.

A protocol is said to achieve a *throughput* λ if, when it is run with arrival rate λ , the average rate of successfully sent messages is λ . The *delay* experienced by a message is the number of steps from its arrival to its successful transmission.

The *capacity* of a protocol is defined as the maximum throughput λ such that the expected delay remains finite for all messages [43]. The capacity can be found by analysing the underlying Markov chain representing the evolution of the protocol.

Molle and Polyzos note that (according to their definition) the capacity is different from the maximum throughput because of the finite expected delay requirement. They present a simple example of a round robin protocol which achieves for the infinite model a maximum throughput of 1, but has capacity 0 since the expected delay is infinite for all $\lambda > 0$ [43, page 17].

Whenever possible, we will not be satisfied solely with showing stability of a protocol. A major concern of ours is to find bounds on the expected load in the system, since a protocol may be positive recurrent and yet may have a very large average load. However, as we will see later, this is not always easy to show, especially for some popular cases such as backoff protocols.

1.4 Related Results

We next review some of the previous results regarding the stability and performance of contention resolution protocols. These results are classified according to the two basic models: the finite and the infinite. We will restrict the scope of this review to the standard slotted model, where time is divided into equal discrete steps or slots.

1.4.1 The Infinite Model

There is a rich body of research in this model dating back to the 70s after the introduction of the ALOHA protocol by Abramson [1]. Active users in the ALOHA

protocol send with constant probability in every step, and thus have no means of adapting to congestion or bursty situations. ALOHA is clearly a simple acknowledgement-based protocol, but unfortunately there exist many results showing that the ALOHA protocol itself is unstable in the infinite model. A number of these results show that the protocol is not positive recurrent for any $\lambda > 0$, and thus, starting from the empty state, the expected return time is infinite [13, 14, 33]. A stronger result showing that the Markov chain representing ALOHA is transient for any $\lambda > 0$ was first proved by Rosenkrantz and Towsley using a martingale approach [48]. A special case of a result by Kelly [34] also shows that the total number of successful transmissions over time in the ALOHA protocol is finite with probability 1.

Researchers also investigated the behaviour of the ALOHA protocol before it enters into an overloaded state, and found that the protocol had a bistable behaviour where it appears to be stable until it reaches a certain state which made it unstable (in the sense that it never returned to the empty state) thereafter [9]. This behaviour was analysed by Drmota and Schmid [12] who calculate the expected length of time the ALOHA protocol works successfully (when it starts from the empty state) before it diverges and becomes transient.

After the clear evidence showing the inherent instability of constant back-off protocols (i.e. the ALOHA protocol), researchers turned their attention to binary exponential backoff. This was mainly due to the popularity of this protocol which is the basis for the Ethernet [40]. Instead of sending with a fixed probability in every step, binary exponential backoff multiplies the probability by half after every collision. Rosenkrantz [47] showed that this protocol is unstable in the sense that the expected backlog (i.e. number of messages waiting to be sent) is infinite whenever $\lambda > 0.72$. Later, Aldous [4] used an intricate proof to show that binary exponential backoff is unstable for any λ in the infinite model. His proof shows that the Markov chain is transient and that the number of successful transmissions is $o(t)$ for any large enough period $0, \dots, t$.

Another type of acknowledgement-based protocols, called age-based proto-

cols were studied in the PhD thesis of MacPhee [38] and Kelly and MacPhee [35]. The authors calculated a critical value λ_c of the arrival rate, and showed that the expected number of successful transmissions is infinite if and only if the arrival rate $\lambda \leq \lambda_c$. Thus when $\lambda > \lambda_c$, the number of successful transmissions is finite and the protocol is unstable. For the ALOHA they showed that $\lambda_c = 0$, and for binary exponential backoff $\lambda_c = \ln(2)$. Ingenoso [31] also considered age-based protocols in his Ph.D. thesis and showed that age-based protocols with monotonically decreasing probabilities $p_0, p_1, \dots$ are transient for any $\lambda > 0$.

More recently, Goldberg, Jerrum, Kannan and Paterson [21] examined the stability question for general backoff protocols, and showed that no backoff protocol is recurrent for arrival rate $\lambda > 0.42$. Thus this result implies that backoff protocols provably cannot do as well as full-sending protocols (which have been shown to achieve a throughput of 0.48776). They also showed that no acknowledgement-based protocol is recurrent for $\lambda \geq 0.530045$. This implies that age-based protocols cannot do better than this, since they belong to the acknowledgement-based class.

Goldberg, MacKenzie, Paterson, and Srinivasan [23] introduced a protocol which is strongly stable for any $\lambda < 1/e$, and where the expected waiting time of any message is constant. The protocol is similar to acknowledgement-based protocols in the sense that it requires binary feedback for the users transmissions only. However, it also assumes that the users are synchronised and share a common clock. Another achievement of this protocol is that it is the first infinite model protocol which is provably strongly stable (i.e. $E[T_{ret}] < \infty$ and $E[L_{avg}] < \infty$).

If we allow the user the power of full sensing (i.e. ternary or m -ary feedback and the complete history of all steps), then many relatively efficient protocols have been proposed for the infinite model. However, this puts many constraints on the protocol which may not be practical in many situations. For example, keeping track of the whole history of the activity in the channel means that all

users must start together from the beginning. Thus users cannot be inactive for a certain period of time and must continuously monitor the channel from step 1. Furthermore, in some situations binary feedback may be all that is available to the users. Molle and Polyzos [43] provide a comprehensive and well written survey of all of the full-sensing protocols in the infinite model, their performance, and methods of analysing them. Their survey also includes some limited sensing and binary feedback protocols.

The discovery of these efficient full sensing protocols dates back to 1978, when Capetanakis [8] and Tsybakov and Mikhailov [52] independently proposed protocols for the infinite model which could achieve capacity of up to 0.347. Their discovery was followed by a sequence of protocols with better capacity leading to the protocol discovered by Gallager [18] (and then rediscovered by Tsybakov and Mikhailov [54]) which could achieve 0.487 capacity. A modification of this protocol was then proposed by Mosely and Humblet [44] which became the most efficient protocol for the infinite model (i.e. in terms of capacity). Modifications to this protocol have been proposed but they all lead to minor improvements in capacity (in the seventh decimal digit). Note also that these protocols can be modified to work in the limited sensing model (i.e. users can switch on and off from sensing the network environment). However, these protocols usually require users who join the system to start in an often lengthy synchronisation process. This initialisation period is used by the new user to calculate some state variables needed by the protocol [43].

1.4.2 The Finite Model

Researchers started considering the finite model by examining the performance of the ALOHA protocol with a fixed set of users. Tsybakov and Mikhailov [53] used a stochastic domination argument and generating functions to specify the conditions under which the Markov chain is either positive recurrent or tran-

sient. These conditions depend on both the arrival rate ³ λ and the retransmission probability chosen by the system (which is a constant in the case of ALOHA). They showed that the optimal operation of ALOHA occurs when the number of users n is known and fixed. When this is the case, if every user sends with probability $1/n$ then the protocol is stable (in the sense that the underlying Markov chain is positive recurrent and the expected waiting time is finite) for any $\lambda < 1/e$ and is unstable otherwise. However, this knowledge of n is clearly a strong (and in some situations impossible) assumption because of the following reasons:

1. The protocol is no longer acknowledgement-based.
2. The protocol is no longer truly distributed since there must be a central controller which determines the value of n and conveys it to the users.
3. The protocol is no longer dynamic (or adaptive), in the sense that it cannot adapt to different population sizes over time. This is important in many applications where the number of users varies over time (such as mobile networks, or Internet applications).

Binary exponential backoff in the finite model was first studied by Goodman, Greenberg, Madras, and March [25] who proved that the protocol is positive recurrent for arrival rates smaller than $n^{-\Theta(\ln n)}$. Al-Ammal, Goldberg, and MacKenzie [3] recently showed (in a result which is generalised in this thesis) that the binary exponential backoff protocol is stable (i.e. the chain is positive recurrent) for arrival rates $\lambda \leq 1/(\alpha n^{0.75+\epsilon})$ for some constant α , any $\epsilon > 0$ and sufficiently large n . On the instability front, Håstad, Leighton, and Rogoff [27] showed that binary exponential backoff is unstable (i.e. both $E[T_{ret}] = \infty$ and $E[L_{avg}] = \infty$) when $\lambda > \lambda_0 + \frac{1}{4n-2}$, where λ_0 is the solution to $\lambda_0 = e^{-\lambda_0}$ and is approximately 0.567. In Chapter 4, we will generalise this result for any c -ary exponential backoff protocol (i.e. exponential backoff with base $c > 1$). Håstad

³Because there are n users in the finite model, most authors divide the total arrival rate λ symmetrically into λ/n for all n users. However, some results allow more flexibility and prove stability or instability for non-symmetric arrival rates [25].

et al. [27] also showed that when n is sufficiently big, then binary exponential backoff is unstable for $\lambda > 1/2$.

Håstad, Leighton, and Rogoff [27] also examined other backoff protocols in the finite model. Their most significant result was for polynomial backoff. They showed that polynomial backoff is a better choice than exponential backoff when it comes to stability. In particular they showed the following:

1. Every polynomial backoff with probability $p_i = (i + 1)^{-\alpha}$ where $\alpha > 1$ is a constant, is stable (i.e. both $E[T_{ret}] < \infty$ and $E[L_{avg}] < \infty$) for any $\lambda < 1$.
2. Every polynomial backoff with probability $p_i = (i + 1)^{-\alpha}$ where $\alpha \leq 1$ is a positive constant, is unstable (in the sense that both $E[T_{ret}]$ and $E[L_{avg}]$ are infinite) for any $\lambda < 1$.

This result produced the first acknowledgement-based protocol that is stable in such a strong sense for such a high arrival rate.

Motivated by the success with polynomial backoff, Goldberg and MacKenzie [22] studied the polynomial backoff protocol when n clients are communicating with k servers, and each server deals with contention like a multiple-access channel. Suppose that at each step client i generates a message for server j with probability $\lambda_{i,j}$, and define the mean client-server arrival rate as the maximum over all clients i and servers j of the sum of the mean arrival rates associated with client i or server j . Goldberg and MacKenzie showed that for any mean client to server arrival rate less than one, polynomial backoff is stable. Their proof used the same strong definition of stability introduced by Håstad, Leighton and Rogoff.

Although the results for polynomial backoff were all optimal in terms of arrival rates and stability, no polynomial bound on the expected waiting time or the expected average load was obtained. Raghavan and Upfal [46] however considered the contention resolution problem for an almost-acknowledgement-based model. They assumed that the users had an estimate of the logarithm of n , the total number of users. They showed that if the arrival rate is less than $1/10$, then the expected waiting time of each message is $O(\log(n))$. This result

also applies to systems with multiple-servers. Goldberg, MacKenzie, Paterson, and Srinivasan [23] designed another contention resolution protocol which also assumes knowledge of an upper bound on n and which is stable (in the strong sense defined by Håstad et al., i.e. where $E[T_{ret}] < \infty$ and $E[L_{avg}] < \infty$) for $\lambda < 1/e$. However, the major achievement of this protocol is that the expected average waiting time of messages is $O(1)$.

1.5 Which protocols are we interested in?

In the preceding review, we have seen that there exists many good protocols which are stable when we can afford one or more of the following:

1. full sensing (including knowledge of the whole past history from step 1 by all users),
2. ternary feedback,
3. some knowledge of the number of users n .

However, in many practical situations these assumptions are not true. We are interested in protocols that are strictly acknowledgement-based, and where the value of n is not known to the users (or where n may vary over a vast range of numbers over time). Consider for example the following situation illustrated in Figure 1.1.

A server is active and is waiting for requests from clients. The number of clients may vary over time, so n is not known. This is typical of Internet-type conditions, where the number of users accessing a server can vary according to the time of day or other factors. Also, the clients have no means of knowing whether the server is busy when they do not send any requests. A request is either answered by an acknowledgement and the requested data, or a busy signal.

This behaviour is typical of TCP/IP network protocols [32, 50]. Although TCP/IP contains a wide range of protocols, those concerned with wide area networking use acknowledgement feedback to limit network traffic. Enabling all




Figure 1.1: An example of a client server protocol on a wide area network such as the Internet. An unknown number of clients competing for a single server on a network with unknown topology. Requests and acknowledgements may be lost so the only indication of success is the reception of an acknowledgement signal from the server. Other clients may join the contention process for the server or may leave without notifying the server or the other clients.

clients in the TCP/IP network to monitor the server's activity even when they are not accessing this server will increase the network traffic by an unacceptable amount. Furthermore, similar to TCP/IP, the network has an unknown topology and it is not possible to calculate the number of clients using the server.

This situation can obviously be modelled by a multiple-access channel representing the server, and users representing the clients. The channel has binary feedback and the protocol controlling the clients *must* be acknowledgement-based. In fact, this situation is somewhat similar to the http protocol which requires that http clients use a modified version of binary exponential backoff [16]. A quick search through the literature shows that binary exponential backoff seems to be very popular among protocol and network designers. Other networking protocols which implement versions of binary exponential backoff include Ethernet and TCP/IP ⁴ [30].

In this thesis we are interested in the study of acknowledgement-based protocols. The reasons for limiting this study to this class of protocols are the following:

1. The class of acknowledgement-based protocols makes very few assumptions about the user's sensing power and the information available to the users. Thus these protocols work in situations where the number of users is unknown or may vary with time, and do not require full sensing. A user can become active or inactive at any stage in the running of the protocol. This simplicity is necessary in many practical situations and for many kinds of networks where users are truly distributed.

⁴Actually, binary exponential backoff has always been a favourite among programmers and network designers mainly because of its simplicity which it derives from being an acknowledgement-based protocol. It requires very little information about the network and the actions of the server. Furthermore, the client or user only needs to keep track of a single integer counter. The admiration for this protocol was stated by one of the designers of the TCP/IP protocol in the following remark [32]: "For a transport endpoint embedded in a network of unknown topology and with unknown, unknowable and constantly changing population of competing conversations, only one scheme has any hope of working – exponential backoff." The author does not present proof or (even empirical) justification for this remark other than Kelly's [34] proof that anything slower than exponential backoff is unstable in the infinite model.

2. Acknowledgement based protocols are not understood as well as other classes such as the tree or stack protocols proposed by Capetanakis[8] and Tsybakov and Mikhailov [53]. In the finite model, many open questions remain about the existence of an acknowledgement-based protocol that is both stable for high arrival rates and has small expected load. Although polynomial backoff is stable for any $\lambda < 1$, the only known bound on its average expected load is exponential in n ⁵.

Furthermore, we have chosen to concentrate on understanding the stability of exponential backoff protocols because of their widespread use in applications and their simplicity.

1.6 Summary of Results

Our main tools for proving that a protocol is stable are Markov chains and an associated nonnegative function called the potential function. In the next chapter, we start by presenting some notation and definitions of stability and related concepts. Then we present some theorems and lemmas for establishing stability of a Markov chain using suitable potential functions. The main contribution of Chapter 2 is a lemma which shows that a Markov chain is stable in a strong sense (i.e. $E[T_{ret}] < \infty$ and $E_\pi[L(X_t)] < \infty$) if the expected drop in the square of the potential is at least equal to the potential that we started with. This drop in the square of the potential is calculated over a state dependent number of steps (which is bounded by a constant).

Although such a lemma exists in the paper by Håstad et al. [27], the authors made a small error which breaks down their proof ⁶. Goldberg and MacKenzie [22] then produced a correct proof for some similar conditions which

⁵Håstad et al. indicate that their analysis of polynomial backoff “reveals that $E[L_{avg}]$ is at most $P((1 - \lambda)^{-1})2^{Q(n)}$ where P and Q are polynomial functions.”

⁶In the proof of their Lemma 3.3, Håstad, Leighton, and Rogoff [27] use a hitting time (which is at least 0) instead of a first return time (which is at least 1). This renders the proof incorrect and unusable, since their analysis depends on their use of a hitting time, and the Theorem they use from Meyn and Tweedie [41] only works with return times.

do not exactly match the requirements of our applications. We develop in Chapter 2 a new proof for this lemma and show that these conditions are sufficient for a strong sense of stability of the Markov chain (where the chain is positive recurrent and the expected value of the potential function is bounded in the steady state). The new proof uses an embedded Markov chain to show that the original chain is strongly stable.

In Chapter 3 we study the stability of a class of protocols which includes binary exponential backoff. For any $c > 1$, sufficiently large constant α and any positive constant ϵ , we show that the general class of c -ary exponential backoff protocols are stable (positive recurrent) for symmetric arrival rates and overall arrival rate $\lambda \leq 1/(\alpha n^{0.75+\epsilon})$. As a special case of this result, we improve the previous known bound on the stability (i.e. positive recurrence) of binary exponential backoff, which was $n^{-\alpha \log n}$ [25]. Note that this result originally appeared in Al-Ammal, Goldberg, and MacKenzie [3] for binary exponential backoff, and is generalised in Chapter 3 for c -ary exponential backoff protocols for any constant $c > 1$.

In Chapter 4, we show that the class of c -ary exponential backoff protocols are unstable (for $c \geq 2$), in the sense that the process has no stationary distribution (i.e. $E[T_{ret}] = \infty$), for any arrival rate $\lambda > \lambda_0 + \frac{1}{4n-2}$, where λ_0 is a constant approximately equal to 0.567. Then we present sufficient conditions under which any system running the binary exponential backoff protocol with two users is unstable. These conditions improve the previously known results [27] for the two user case from $\lambda > 0.733667$ to $\lambda > 2/3$ when arrival rates are symmetric. We also give some conditions for non-positivity for general non-symmetric arrival rates.

Next, we turn our attention to age-based protocols in Chapter 5. We present an age-based protocol which is similar to binary exponential backoff, and show that it is strongly stable for $\lambda \leq 1/(\alpha n^{0.75})$ for some constant α and big enough n . Seeking some general results for age-based protocols, we investigate the stability of protocols which send messages aged t steps with probability $p_t = t^{-\epsilon}$,

where $0 < \varepsilon < 1$. Under some weaker definition of stability than used for the first age-based protocol (i.e. positive recurrence only), we show that such protocols are stable as long as $\lambda \leq 1/(\alpha n^{1+1/\varepsilon})$. We end the thesis with some remarks, a discussion of other methods for showing stability, and open questions.

Chapter 2

Stability Conditions for Markov Chains

Markov chains are simple processes which seem to capture the behaviour and structure of many real life applications and phenomena. These applications span a wide range of disciplines including control, queueing theory, financial modelling, biology, telecommunications, and computer science among others. The present chapter is mainly concerned with presenting results about Markov chains that will be useful later as tools for showing the stability of contention resolution protocols. These theorems specify conditions necessary for a Markov chain to be stable (under different conditions of stability).

Some of the ideas and methods developed in this chapter have been adopted from Håstad, Leighton, and Rogoff [27]. However, as was stated in the summary in Chapter 1, the main lemma used by them to show stability when certain drift conditions are true (Lemma 3.3 in [27]) has an incorrect proof. Also, some of the details of their use of a result by Meyn and Tweedie are missing. We develop a new proof for this lemma and fill in the missing details for using the Meyn and Tweedie theorem.

2.1 Definitions and Notation

We begin with some notation concerning Markov chains and their properties. Most of this is standard notation that can be found in any good text in applied probability or stochastic processes [26]. However, we restate these definitions here for convenience and completeness. Note that the notation is very similar to that used by Tweedie and Meyn [53]. The reader is referred to their book and to Grimmit and Stirzaker [26] for more information.

A *Markov chain* is a sequence of random variables $\{X_0, X_1, \dots\}$, where the random variable X_t is defined on some countable state space $\mathbf{X}$. The basic feature of a Markov chain, as opposed to other sequences of random variables, is that it adheres to the following property.

Definition 2.1 *The process X_t is a Markov chain if it satisfies the Markov property:*

$$\Pr(X_t = s \mid X_0 = x_0, \dots, X_{t-1} = x_{t-1}) = \Pr(X_t = s \mid X_{t-1} = x_{t-1})$$

for all $t \geq 1$ and all $s, x_0, x_1, \dots, x_{t-1} \in \mathbf{X}$.

Informally, this property indicates that the future actions in a Markov chain depend on its current state only, and not on how it arrived to this state.

The evolution of the chain is determined by its *transition probabilities*. Let $p(x, y)$ denote the single step transition probability from state x to state y . All of the chains described in this thesis are *time-homogeneous*, and thus have the following property for all t ,

$$\Pr(X_{t+1} = x \mid X_t = y) = \Pr(X_1 = x \mid X_0 = y).$$

This property ensures that the transition probabilities do not depend on t , but on the specific state the chain is in. Thus, in homogeneous chains $p(x, y) = \Pr(X_t = y \mid X_{t-1} = x)$ for all t .

Next, we present some more notation regarding the transition probabilities of the chain. Let

$$p(x, A) = \sum_{y \in A} p(x, y)$$

be the one step transition probability from state x to any state in the set A . The i -step transition probability

$$p^i(x, y) = \Pr(X_t = y \mid X_{t-i} = x)$$

is the probability that the chain will move from state x to state y in exactly i steps. Similarly, the probability that the chain, starting in state x , will be in any state in set A in exactly i steps is defined as

$$p^i(x, A) = \sum_{y \in A} p^i(x, y).$$

The *first return time* to a set A is defined as

$$\tau_A = \min\{t \geq 1 : X_t \in A\}.$$

Throughout this chapter the symbols x, y , and z will be used to denote states in $\mathbf{X}$, while A, B , and C will denote subsets in $\mathbf{X}$. We denote by $\Pr_x(E)$ the probability of event E conditioned on $X_0 = x$. The probability that the chain *ever enters* set A starting from state x is denoted by

$$L(x, A) = \Pr_x(\tau_A < \infty).$$

Next we present definitions for the standard properties of aperiodicity and irreducibility, and then define the stationary distribution of a Markov chain.

Definition 2.2 (Aperiodicity) *The period $d(x)$ of a state x is defined by*

$$d(x) = \gcd\{i : p^i(x, x) > 0\}.$$

We call x periodic if $d(x) > 1$ and aperiodic if $d(x) = 1$. A chain is called aperiodic if all of its states are aperiodic.

Definition 2.3 (Irreducibility) *The chain X_t is called irreducible if for any pair x, y of states we have $p^m(x, y) > 0$ for some $m \geq 0$.*

A Markov chain which is both irreducible and aperiodic may have a stationary distribution. This distribution is a limiting distribution which is independent of the starting state of the Markov chain, and specifies the probabilities of being in a certain state once the chain enters this steady state. Basically, this distribution is a good description of the long term behaviour of the chain.

Definition 2.4 (Stationary distribution) *The vector π is called a stationary distribution of the chain X_t if π has components $(\pi_x : x \in \mathbf{X})$ such that*

- (a) $\pi_x \geq 0$ for all x , and $\sum_x \pi_x = 1$.
- (b) $\pi_x = \sum_y \pi_y p(x, y)$ for all x .

It is a well known fact in probability theory that, when such a distribution exists, and if the Markov chain is in the stationary distribution at step t , it remains in the stationary distribution at step $t + 1$. Intuitively, the stationary distribution describes the steady-state behaviour of the Markov chain. This distribution is also a limiting distribution, and the chain approaches it as time approaches infinity. Although all Markov chains with a finite state space possess a stationary distribution, in the countable state space case we must explicitly show its existence. Thus, throughout this thesis our goal will be to prove the existence (or nonexistence) of this stationary distribution for various denumerable Markov chains modelling contention resolution protocols. We will also attempt to characterise some properties of the system, such as the system load, when the chain enters this stationary distribution.

2.2 Definitions of Stability

The main goal of this thesis is the study of the performance of contention resolution protocols. The most obvious and intuitive measure for the success of a protocol is its ability to deliver the incoming packets to their destination in a reasonable amount of time. This can be measured in several ways. For example, we can calculate the average waiting time W_{avg} of a message before it is

successfully delivered to its destination. Practitioners and engineers might be interested in the average number of waiting messages over time L_{avg} to decide on appropriate queue sizes at each station.

Even though these two measurements might seem different at first sight, previous research has shown that they are actually closely related. In fact, it was shown that with probability one,

$$L_{avg} = \lambda W_{avg}$$

where λ is the overall arrival rate of messages into the system [51]. Since the protocols that will be studied here are all randomised, it is crucial that the expected average load $E[L_{avg}]$ and the expected average waiting time $E[W_{avg}]$ be as small as possible. Clearly, we would like to prove that at least $E[L_{avg}] < \infty$.

Another notion of stability, which is very popular in the applied probability literature requires that the chain returns to the starting state with probability one. This property is called *recurrence*. A stronger stability property, called *positive recurrence* requires that $E[T_{ret}] < \infty$. Both properties are formally defined as follows.

Definition 2.5 *A chain is said to be recurrent if it returns to the start state with probability one. The chain is said to be positive recurrent if $E[T_{ret}] < \infty$.*

Definition 2.6 *A contention resolution protocol is said to be positive recurrent if $E[T_{ret}] < \infty$.*

We will mostly be concerned with positive recurrence since it is a stronger property than recurrence and it guarantees the existence of a unique stationary distribution in all irreducible homogeneous Markov chains [7, page 104].

For the contention resolution protocols that will be studied later, this is equivalent to showing that the expected time needed for the system to empty its queues is finite. Thus assuming that the contention resolution protocol can be modelled by a Markov chain, we can use tools from probability theory to

show that $E[T_{ret}] < \infty$. These tools include potential (or Lyapunov) functions and drift conditions which will be described in the next section.

Ultimately, we are interested in studying the stability of contention resolution protocols. Thus we will next define a stronger stability concept than positive recurrence for a system of n users running such a protocol. We will assume that each user holds a queue of messages waiting to be sent. To cover both performance measures described above, we will require that both $E[T_{ret}]$ and $E_\pi[L(X_t)]$ be finite in our definition of stability.

Now we are ready to present the formal definition for stability of a contention resolution protocol.

Definition 2.7 *A contention resolution protocol is said to be strongly stable if both of the following conditions are satisfied*

1. $E[T_{ret}] < \infty$, and
2. $E_\pi[L(X_t)] < \infty$.

This is equivalent to saying that the Markov chain X_t modelling the protocol is positive recurrent and that the expected load of the system in the steady state is finite. This definition of stability is a strong one since it requires the existence of a stationary distribution and for the load function to be finite under this distribution. A similar definition was first adopted by Håstad, Leighton, and Rogoff [27] in an attempt to unify the different versions of stability that appear in the contention resolution literature. However, they require positive recurrence and $E[L_{avg}] < \infty$.

Note that if there exists a function f such that $L(X_t) \leq f(X_t)$ and the chain X_t has a stationary distribution π with $E_\pi[f(X_t)] < \infty$, then we can conclude that $E_\pi[L(X_t)] < \infty$. Therefore, we will often pick a potential function f which is an upper bound on the system load, and show that the expected value of this potential function is finite in the steady state.

Let $\mathcal{B}(\mathbf{X})$ be the power set of the state space $\mathbf{X}$. The following definition will be necessary later.

Definition 2.8 A set $C \subseteq \mathbf{X}$ is called *small* if there exists an $m > 0$ and a non-trivial measure ν_m on $\mathcal{B}(\mathbf{X})$, such that for all $x \in C$, and $B \subseteq \mathbf{X}$,

$$p^m(x, B) \geq \nu_m(B). \quad (2.1)$$

2.3 Establishing Positive Recurrence and Strong Stability Using Drift

Let X_t , for $t = 0, 1, \dots$ be a Markov chain defined on a countable state space $\mathbf{X}$. The most common tool for proving that a Markov chain is positive recurrent is Foster's theorem [17].

Theorem 2.1 (Foster) *A time-homogeneous irreducible aperiodic Markov chain X_t with a countable state space $\mathbf{X}$ is positive recurrent iff there exists a positive function $f(\rho)$, $\rho \in \mathbf{X}$, a number $\epsilon > 0$, and a finite set $C \subseteq \mathbf{X}$, such that the following inequalities hold.*

$$E[f(X(t+1)) - f(X(t)) \mid X(t) = \rho] \leq -\epsilon, \quad \rho \notin C \quad (2.2)$$

$$E[f(X(t+1)) \mid X(t) = \rho] < \infty, \quad \rho \in C. \quad (2.3)$$

Basically, the idea is to use a “potential¹ function” f to follow the progress of the chain. The function f is positive for all states in $\mathbf{X}$. The theorem states that the chain is positive recurrent iff there exists a potential function f which

1. usually decreases (Equation (2.2)), and
2. cannot increase much (Equation (2.3))

in a single step. Equation (2.2) implies that, from any state $\rho \notin C$, the expected time to reach C from ρ is at most $f(\rho)/\epsilon$. This (combined with Equation (2.3)) implies that the expected return time to C is finite, which in turn implies that the chain is positive recurrent. (For more details on Foster's theorem and its proof see [15, 41].)

¹This function is sometimes called a Lyapunov or test function by other authors.

2.3.1 Multiple step drifts

We use the word *drift* to describe the expected change in the potential function, $E[f(X_{t+1}) - f(X_t)]$. This is sometimes called the *single step drift*, and other multiple step drifts can be defined where the expected change is taken over more than one step.

The original Foster's theorem works by examining single step drifts. In practice, it can be difficult to find a potential function satisfying the criteria in Foster's theorem. Many natural functions will not exhibit negative drift in a single step. However, in some situations, it can be easier to show that the drift is negative when considering multiple steps. For example, consider a protocol running on a multiple-access channel where there are two stations which will send with probability one. A collision in this case is unavoidable, and no improvement in any natural potential function is possible in a single step. However, it is reasonable to assume that this situation will change after a certain number of steps.

Tools for showing positive recurrence when the multiple step drift is negative are available in the applied probability and Markov chain literature [41, 15]. For example, the following generalisation of Foster's theorem [15] will be used later to show positive recurrence of some contention resolution protocols.

Theorem 2.2 (Foster; Fayolle, Malyshev, Menshikov)

A time-homogeneous irreducible aperiodic Markov chain X with a countable state space $\mathbf{X}$ is positive recurrent iff there exists a positive function $f(s)$, $s \in \mathbf{X}$, a number $\epsilon > 0$, a positive integer-valued function $k(s)$, $s \in \mathbf{X}$, and a finite set $C \subseteq \mathbf{X}$, such that the following inequalities hold.

$$E[f(X_{t+k(X_t)}) - f(X_t) \mid X_t = s] \leq -\epsilon k(s), \quad s \notin C \quad (2.4)$$

$$E[f(X_{t+k(X_t)}) \mid X_t = s] < \infty, \quad s \in C. \quad (2.5)$$

2.3.2 Strong stability

In this section we will derive some criteria for showing *strong stability* with multiple step drift. As was stated earlier, this kind of stability requires positive recurrence of the chain and a bound on the expected load in the steady state. We will achieve this using an embedded process which jumps over these multiple steps. This process is also a Markov chain using the strong Markov property, and it will enable us to show stability for the original chain.

We start by presenting the following definition of φ -irreducibility which is copied from Meyn and Tweedie [41, page 87].

Definition 2.9 *We call X_t φ -irreducible if there exists a measure φ on $\mathcal{B}(\mathbf{X})$ such that, whenever $\varphi(A) > 0$, we have $L(x, A) > 0$ for all $x \in \mathbf{X}$.*

If there exists a measure φ which satisfies Definition 2.9 then the Markov chain is also ψ -irreducible for a unique *maximal irreducibility measure* ψ using the following proposition [41, page 88].

Proposition 2.3 *If X_t is φ -irreducible for some measure φ , then there exists a probability measure ψ on $\mathcal{B}(\mathbf{X})$ such that*

1. X_t is ψ -irreducible;
2. for any other measure φ' , the chain X_t is φ' -irreducible if and only if $\psi > \varphi'$.

We will show now that every countable state space, irreducible and aperiodic chain is φ -irreducible, and thus there exists a maximal irreducibility measure ψ for which the chain is ψ -irreducible and every other irreducibility measure is dominated by ψ . Note that we will not explicitly use the properties of ψ -irreducibility. We are only interested in ensuring that the chain X_t has this property in order to use Theorem 2.4 below to show that X_t is strongly stable.

In general, we can show that every countable state space Markov chain which is irreducible and aperiodic is ψ -irreducible by specifying the measure φ as the *counting* (or cardinality) measure. In other words, for every set $A \subseteq \mathbf{X}$,

$$\varphi(A) = |A|.$$

Clearly, this satisfies Definition 2.9 since any nonempty set A is reachable from any state x in an irreducible countable state space chain. Therefore, we conclude that X_t is ψ -irreducible according to Proposition 2.3. In general, suppose that the chain is irreducible, then it is φ -irreducible for any measure φ which assigns measure 0 to the empty set. Then by Proposition 2.3 there exists a probability measure ψ for which it is ψ -irreducible.

The following theorem will be invoked later to show that (under some conditions) the Markov chain X_t is strongly stable. It is due to Meyn and Tweedie and is Theorem 14.0.1 in their book [41].

Theorem 2.4 (f -Norm Ergodic Theorem)

Suppose that a chain X_t is ψ -irreducible and aperiodic, and let $f \geq 1$ be a function defined on its state space $\mathbf{X}$. Then the following statements are equivalent:

- (i) *The chain is positive recurrent with invariant probability measure π and*

$$\pi(f) := \int \pi(dx) f(x) < \infty.$$

- (ii) *There exists some small set $C \subseteq \mathbf{X}$ such that*

$$\sup_{x \in C} E_x \left[\sum_{t=0}^{\tau_C-1} f(X_t) \right] < \infty. \quad (2.6)$$

- (iii) *There exists a constant b and a small set C and some function $V : \mathbf{X} \rightarrow \mathbb{R}^+ \cup \infty$ satisfying $V(x_0) < \infty$ for some $x_0 \in \mathbf{X}$, and*

$$E[V(X_{t+1}) - V(X_t)] \leq -f(X_t) + b\mathbf{1}_C(X_t). \quad (2.7)$$

The f -Norm Ergodic Theorem is stated in terms of general state space chains. Since the Markov chains we will be studying are all countable state space irreducible and aperiodic chains, we can make some assumptions to simplify the requirements of the theorem.

Note also that there were two changes that we have made while restating Theorem 14.0.1 from Meyn and Tweedie [41] as Theorem 2.4 here. First, note that the original statement of the theorem in Meyn and Tweedie uses petite sets

in describing set C in statements (ii) and (iii). We have replaced the word petite by small because every small set is petite² (see Meyn and Tweedie [41], p. 121 and Proposition 5.5.3). So, if a small set C exists with the desired properties, then it can be considered as petite and may be used in the theorem. Secondly, we must show that the chain X_t is ψ -irreducible.

Because of the countable state space and the properties of irreducibility and aperiodicity of X_t , we can show that every finite subset of a countable state space is a small set. This will be useful later, since we will set C in Theorem 2.4 to a finite set.

Proposition 2.5 *In an irreducible and aperiodic Markov chain with countable state space $\mathbf{X}$, every finite set C , where $C \subseteq \mathbf{X}$, is small.*

Proof: To show that every finite set C of states is small, we must show that there exists a non-trivial measure ν_m on $B(\mathbf{X})$ such that for all $x \in C$, and $B \subseteq \mathbf{X}$,

$$p^m(x, B) \geq \nu_m(B).$$

Since X_t is irreducible, by definition, for any two states x and y

$$p^i(x, y) > 0 \quad \text{for some } i \geq 0.$$

Also, since X_t is aperiodic, then for any two states x and y

$$p^i(x, y) > 0 \quad \text{for a sufficiently large } i.$$

Let us set $\nu_m(B) = \delta \varphi(B)$ for some $\delta > 0$ and for $\varphi(B) = |B|$ the counting measure. Then for all $x \in C$,

$$p^m(x, B) = \sum_{y \in B} p^m(x, y) > 0 \quad \text{for any nonempty set } B \subseteq \mathbf{X}.$$

This is true since irreducibility and aperiodicity guarantee the existence of a sufficiently large m such that this is true. This proves the existence of the non-trivial measure $\nu_m(B)$, and therefore C is small by Definition 2.8. $\square$

²Actually, the converse is also true when the Markov chain is irreducible and aperiodic. See Theorem 5.5.7 in [41].

Next we present a lemma with necessary conditions for strong stability. It specifies two conditions which are sufficient to show that a Markov chain is positive recurrent and that the potential function has a finite expected value in the stationary distribution. This lemma will enable us to show stability by examining the drift over multiple steps instead of the single step criterion used in Theorem 2.4. Furthermore, the number of steps that we will examine depends on the starting state. This will give us more flexibility when dealing with different kinds of starting situations depending on the starting state (such as big queues or large backoff counters). As we pointed out earlier, this lemma was suggested by Håstad et al. [27]. We present here a completely new proof for it which uses an embedded chain and the f -norm ergodic theorem.

Lemma 2.6 *Suppose that a chain X_t has a countable state space and is irreducible and aperiodic, and let $f \geq 1$ be a function defined on the state space $\mathbf{X}$. If there exists a constant $K < \infty$ and a function $k : \mathbf{X} \rightarrow \{1, \dots, K\}$, and both of the following conditions are true.*

(V1) *There exists a constant $b < \infty$ and a finite set $C \subseteq \mathbf{X}$ and some non-negative function $V : \mathbf{X} \rightarrow \mathbb{R}^+ \cup \infty$ satisfying $V(x_0) < \infty$ for some $x_0 \in \mathbf{X}$, and*

$$E[V(X_{t+k(X_t)})] - V(X_t) \leq -f(X_t) + b1_C(X_t).$$

(V2) *There exists a constant d such that for every starting state X_t ,*

$$f(X_{t+1}) < d f(X_t).$$

Then the chain X_t is positive recurrent with invariant probability measure π and $E_\pi[f(X_t)] < \infty$.

Proof:

The main difficulty in applying Theorem 2.4 to the proof of this lemma is that condition (V1) is true over (possibly) multiple steps while statement (ii) in the theorem uses a single step only. To deal with this we will work initially with an embedded Markov chain $\hat{X}_t$. The same construction of $\hat{X}_t$ has been used by

Meyn and Tweedie for obtaining state dependent drift criteria (see Meyn and Tweedie [41], page 466). Using the function $k(x)$ we will define a new transition law for $\widehat{X}_t$ as follows

$$\widehat{p}(x, y) = p^{k(x)}(x, y),$$

where $p^i(x, y)$ is the probability that starting from state x the chain will end up in state y in exactly i steps.

The embedded Markov chain can be constructed as follows. First, note that $k(x)$ as defined in the hypothesis of the lemma is a stopping time. Let $s(i)$ denote the iterates of this stopping time. In other words, let

$$\begin{aligned} s(0) &= 0, \\ s(1) &= k(X_0), \\ &\vdots \\ s(i+1) &= s(i) + k(X_{s(i)}). \end{aligned}$$

Using the strong Markov property (see for example Brémaud [7] page 85) we conclude that

$$\widehat{X}_i = X_{s(i)}, \quad i \geq 0$$

is also a Markov chain with transition probabilities $\widehat{p}$. Informally, the embedded chain will be jumping along some of the states in any path of the original chain. Its states will coincide with the original chain's states at exactly the stopping times defined above. For an illustration of the embedded chain and its relation to the original chain, refer to Figure 2.1.

Notice that we are using a finite set C in the statement of this lemma, instead of the small set specified in Theorem 2.4. This is allowed since in the countable state space setting when the chain is irreducible and aperiodic we can show that every finite set $C \subseteq \mathbf{X}$ is small. This has been established using Proposition 2.5.

Let $\widehat{\tau}_C$ be the first return time by the embedded chain $\widehat{X}_t$ to the set C . Notice that $s(\widehat{\tau}_C)$ denotes the return time by the original chain X_t along an




Figure 2.1: A sample path of the embedded chain. The states of the two chains coincide on the stopping times.

embedded path, and it is easy to see that

$$\tau_C \leq s(\hat{\tau}_C). \quad (2.8)$$

Notice now that using condition (VI) for the embedded chain and in a single step,

$$E[V(\hat{X}_{t+1}) - V(\hat{X}_t)] \leq -f(\hat{X}_t) + b\mathbf{1}_C(\hat{X}_t).$$

Therefore, statement (iii) in Theorem 2.4 is satisfied by the embedded chain. This leads us to conclude that $\hat{X}_t$ is positive recurrent with stationary distribution $\hat{\pi}$ and that $E_{\hat{\pi}}[f(x)] < \infty$.

Since the three statements in Theorem 2.4 are equivalent, from statement (ii) we conclude that for the embedded chain $\hat{X}_t$, there exists a small set $C \subseteq \mathbf{X}$ such that

$$\sup_{x \in C} E\left[\sum_{i=0}^{\hat{\tau}_C-1} f(\hat{X}_i)\right] < \infty. \quad (2.9)$$

Returning back to the original chain X_t , starting from any state $X_0 \in C$,

$$\begin{aligned}
E\left[\sum_{i=0}^{\tau_C-1} f(X_i)\right] &\leq E\left[\sum_{i=0}^{s(\hat{\tau}_C)-1} f(X_i)\right] \\
&= E\left[\sum_{i=0}^{\hat{\tau}_C-1} \sum_{j=0}^{k(\hat{X}_i)-1} f(X_{s(i)+j})\right] \\
&\leq E\left[\sum_{i=0}^{\hat{\tau}_C-1} \sum_{j=0}^{K-1} f(X_{s(i)+j})\right] \\
&\leq E\left[\sum_{i=0}^{\hat{\tau}_C-1} \sum_{j=0}^{K-1} d^j f(\hat{X}_i)\right] \\
&= E\left[\sum_{i=0}^{\hat{\tau}_C-1} f(\hat{X}_i) \sum_{j=0}^{K-1} d^j\right] \\
&= \sum_{j=0}^{K-1} d^j E\left[\sum_{i=0}^{\hat{\tau}_C-1} f(\hat{X}_i)\right] \\
&< \infty.
\end{aligned}$$

The first inequality is obtained using Inequality (2.8). The first equality is true since we can decompose the summation into the sum of the values of the function f at the stopping times (i.e. in the embedded chain) and the sum of the values of f at the states between these stopping times in the original chain. The second inequality is true since the function k is bounded from above by the constant K , and the third inequality is obtained from condition (V2). Finally, we apply the bound in Inequality (2.9) to show that the whole thing is finite.

Notice that we have now proved that for the original chain, there exists a small set C such that

$$\sup_{x \in C} E\left[\sum_{i=0}^{\tau_C-1} f(X_i)\right] < \infty.$$

Using Theorem 2.4 again (but this time on the original chain X_t), we conclude that statement (ii) is true. Hence, using statement (i), the original chain X_t is positive recurrent with stationary distribution π and the expected value of the function f in the stationary distribution is finite. $\square$

Although for most systems it is easy to check if condition (V2) is satisfied

by the underlying Markov chain and the specified potential function f , finding a suitable function V that will satisfy condition (V1) is usually hard. Therefore, we will show in the next lemma that when we set V to be the square of the function f , condition (V1) can be replaced by two conditions which will be simpler to verify for the systems that will be studied later. Note that this lemma is due to Håstad, Leighton, and Rogoff [27] who used it in proving that polynomial backoff is stable.

Lemma 2.7 *Suppose that a chain X_t has a countable state space and is irreducible and aperiodic, and let $f \geq 1$ be a function defined on its state space $\mathbf{X}$. Suppose that there exists a constant K and a function $k : \mathbf{X} \rightarrow \{1, \dots, K\}$. Let us define the function $V(x) = f^2(x)$. If C is a finite set, σ is some positive constant, and for every state X_t in $\mathbf{X} - C$ the following conditions are true:*

$$(C1) \quad E[(f(X_{t+k(X_t)}) - f(X_t))^2] \leq (2\sigma - 1) k(X_t) \cdot f(X_t),$$

$$(C2) \quad E[f(X_{t+k(X_t)}) - f(X_t)] \leq -\sigma k(X_t),$$

and condition (V2) is also true, then condition (V1) is satisfied.

Proof: For convenience, define $\tilde{d} = f(X_{t+k(X_t)}) - f(X_t)$. If we assume that $X_t \in \mathbf{X} - C$, then from the assumptions of the lemma, $E[\tilde{d}] \leq -\sigma k(X_t)$ and $E[\tilde{d}^2] \leq (2\sigma - 1) k(X_t) \cdot f(X_t)$. These bounds are used to calculate

$$\begin{aligned} E[V(X_{t+k(X_t)}) - V(X_t)] &= E[f^2(X_{t+k(X_t)}) - f^2(X_t)] \\ &= E[(f(X_t) + \tilde{d})^2 - f^2(X_t)] \\ &= E[2\tilde{d} \cdot f(X_t) + \tilde{d}^2] \\ &= 2f(X_t) \cdot E[\tilde{d}] + E[\tilde{d}^2] \\ &\leq 2f(X_t) \cdot (-\sigma k(X_t)) + (2\sigma - 1) k(X_t) \cdot f(X_t) \\ &= -f(X_t)k(X_t) \\ &\leq -f(X_t) \end{aligned}$$

since $k(X_t) \geq 1$.

Since C is finite, there exists b such that when $X_t \in C$ then

$$E[V(X_{t+k}) - V(X_t)] \leq d^K V(X_t) \leq -f(X_t) + b.$$

In particular, we can choose

$$b = \max_{x \in C} d^K V(x) + f(x) < \infty$$

since C is a finite set and K is a constant. This shows that condition (V1) is satisfied. $\square$

For convenience, we will restate the three new conditions that are sufficient for the stability of the Markov chain in the following lemma.

Lemma 2.8 *Suppose that a chain X_t has a countable state space and is irreducible and aperiodic, and let $f \geq 1$ be a potential function defined on the state space $\mathbf{X}$, with $L(x) < f(x)$ for all $x \in \mathbf{X}$. Let C be a finite subset of $\mathbf{X}$. If there exists a constant $K < \infty$ and a function $k : \mathbf{X} \rightarrow \{1, \dots, K\}$, and all of the following conditions are true:*

(C1) *If $X_t \in \mathbf{X} - C$ then*

$$E[(f(X_{t+k(X_t)}) - f(X_t))^2] \leq (2\sigma - 1) k(X_t) \cdot f(X_t).$$

(C2) *If $X_t \in \mathbf{X} - C$ then*

$$E[f(X_{t+k(X_t)}) - f(X_t)] \leq -\sigma k(X_t).$$

(V2) *There exists a constant $d < \infty$ such that, for every starting state X_t ,*

$$f(X_{t+1}) < d f(X_t).$$

Then the chain X_t is positive recurrent with invariant probability measure π where $E[T_{ret}] < \infty$ and $E_\pi[L(X_t)] < \infty$.

Proof: From Lemmas 2.6 and 2.7 we can show that when the three conditions are satisfied then the chain X_t is positive recurrent with stationary distribution π and that $\lim_{t \rightarrow \infty} E[f(X_t)] = E_\pi[f(X_t)] \leq M$, for some $M < \infty$. This leads us to conclude that since for all states $x \in \mathbf{X}$, $L(x) \leq f(x)$, then $E_\pi[L(X_t)] < \infty$. $\square$

Chapter 3

The Stability of c -ary Exponential Backoff

The stability question for polynomial backoff protocols in the finite users model was completely answered by Håstad, Leighton, and Rogoff [27]. They showed that these protocols are stable when the backoff function is a superlinear polynomial function, and are unstable for linear or sublinear functions. The class of exponential backoff protocols, however, still has many open questions regarding their stability. The popular Binary exponential backoff protocol received some attention when it was studied by Goodman, Greenberg, Madras, and March [25] who showed that the protocol is positive recurrent for $\lambda < n^{-\Theta(\ln(n))}$. Håstad, Leighton and Rogoff [27], on the other hand, showed that the protocol is unstable for $\lambda > 1/2$ for sufficiently large n .

In this chapter, we study the general class of exponential backoff protocols with base c , for any constant $c > 1$. We show that these protocols are all positive recurrent for $\lambda \leq \Theta(n^{-(0.75+\epsilon)})$ for any constant $\epsilon > 0$. A special case of this result improves the results of Goodman et al. [25] for binary exponential backoff. We defer the question of instability for c -ary exponential backoff to Chapter 4.

Note that a version of the proof for the stability result appeared in a joint paper by Al-Ammal, Goldberg, and MacKenzie [3], where the authors showed

that binary exponential backoff is positive recurrent for the same arrival rate. Here, we generalise this result to show that c -ary backoff is stable (i.e. positive recurrent) for the same arrival rate.

3.1 The Model and the Protocol

Consider a system of n users sharing a single multiple-access channel. As was stated in Chapter 1, the channel is not centrally controlled, and there is no other means of communication or coordination among the users. A user has no means of knowing the total number of users in the system. Each user maintains a queue of the messages it wishes to send. User i also maintains a queue length counter q_i and a backoff counter b_i . Time is divided into discrete time steps. Immediately before the beginning of the t 'th time step, the length of the queue of the i 'th user is denoted by $q_i(t)$, and the number of times that the message at the head of its queue has collided while trying to send is denoted by $b_i(t)$.

At the beginning of the t 'th step, user i receives independently one message with probability λ_i , and none with probability $1 - \lambda_i$. The total mean arrival rate for the system is $\lambda = \sum_{i=1}^n \lambda_i$. After the message is added to the end of the queue, each user makes an independent decision about whether or not to send the message at the head of its queue. This decision is controlled by the contention resolution protocol.

The protocol considered here is c -ary exponential backoff which is described as follows. If a user has an empty queue, it stays silent and its backoff counter remains zero. If the queue of user i is not empty and the head of its queue has had $b_i(t)$ collisions so far, then it will send with probability $c^{-b_i(t)}$, where $c > 1$ is a real constant. This means that a message at the head of the queue that has experienced no collisions will be sent immediately and that every collision will decrease the probability of an attempt by a factor of $1/c$.

Since the users are sharing a single multiple access channel, if more than one user send simultaneously at any step, a collision is detected and all the

transmissions fail. On the other hand, if a single user sends at that step the transmission succeeds, the message is removed from the head of the queue, and the backoff counter of that user is reset to zero. Users which send on a certain time step monitor the state of the channel and can identify if the previous transmission attempt was a success or a failure. Any user participating in a collision updates its backoff counter by incrementing it by one.

The evolution of the whole system can be modelled by a Markov chain in which the state just before step t is

$$X_t = (q_1(t), \dots, q_n(t), b_1(t), \dots, b_n(t)).$$

The next state X_{t+1} is determined by the arrival process and the actions of the n users which execute the c -ary exponential backoff protocol described above.

3.2 Choosing a Suitable Potential Function

Before presenting the results of this chapter, let us consider the importance of choosing a suitable potential function f . As we have seen in Section 2.3, Foster's theorem (Theorem 2.1) establishes that if the Markov chain is in fact positive recurrent, then there exists a potential function with negative single step drift. However, finding this potential function can be a highly non-trivial task. In fact, this could be difficult even if we consider the change in the potential over multiple steps.

In the excellent paper by Håstad, Leighton and Rogoff [27], the authors studied the performance of backoff protocols under the same model outlined above. Among other results, they showed that while binary exponential backoff is unstable for $\lambda > 1/2$, polynomial backoff is stable for any arrival rate $\lambda < 1$. The argument used to prove their stability result includes a capture analysis, which shows that a user with a huge queue has a good chance of dominating the channel for a long time. This enables this user to empty its queue and thus creates a very big reduction in the potential. In this section we will argue that using a potential function and case division similar to those used by Håstad

et al. [27] for polynomial backoff it is not possible to show stability results for exponential backoff.

Suppose that we are studying a backoff protocol with a backoff function $h(b)$. (For example, binary exponential backoff has $h(b) = 2^{-b}$.) For attacking stability problems in a backoff system, Håstad, Leighton, and Rogoff suggested using a potential function of the form

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

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

This type of function seems natural since it includes the total system load (i.e. the sum of the sizes of the queues) as well as the sum of inverse powers of the current backoff probabilities. A natural potential function aims to reflect how far is the current state from the empty state where all packets are successfully sent. However, for many systems, using the system load only as a measure of the distance from the start state is not enough. For example, consider a system running binary exponential backoff where only one of the users has a single message and has a backoff counter which is huge. We can expect this message to wait for a very long time because it will have a very low probability of sending. Thus, the second term has been added to the potential function to include such cases. For a user i with backoff counter b_i , the expected number of steps before it will attempt to send is $(h(b_i))^{-1}$. The second term with $c_2 = 1$ can be considered as the sum of the expected waiting times of all users for the first attempt to send.

After specifying a natural potential function, it is vital to find conditions under which this function is expected to decrease in a finite number of steps. This is usually done by partitioning the state space of the Markov chain into cases. In each case, there will be some reason which will drive the expected potential to drop. This coupling between the partition of the state space and the given potential function is another difficult part of the proof. The cases can present the crucial idea behind the argument of the proof.

For example, Håstad, Leighton and Rogoff [27] showed that polynomial backoff with $h(b_i) = (b_i + 1)^{-\alpha}$ for any constant $\alpha > 1$ is stable using the following potential function

$$f(X_t) = C \sum_{i=1}^n q_i + \sum_{i=1}^n (b_i + 1)^{\alpha+1/2} - n. \quad (3.1)$$

They showed that for all states with sufficiently large potential, this function is expected to drop in a finite number of steps. Their proof technique relied on a case partition similar to the following.

Suppose that we are in a state x such that $f(x) > V_0$ for some large constant V_0 . Let B be a very large constant.

- Case 1. $\forall i \ b_i \leq B$.
- Case 2. $\exists i \ b_i > B$.

In case 1, all counters are small. Since they assumed that the potential is at least V_0 for some large constant V_0 , there must be at least one very large queue. Assume without loss of generality that this is queue number n . This queue has at least $q_n \geq C^{-1} \left(\frac{V_0}{n} + (B + 1)^{\alpha+1/2} \right)$ messages. We can now use this information to show that this queue is sufficiently likely to send successfully and dominate the channel by sending until it empties its queue. This will create a very large decrease in the potential. The main argument for this case is a very long and detailed proof that this happens with a not too small probability. They prove this using a capture argument where the authors show that a single user dominates the channel for a very long time. This capture effect is a nice feature of backoff protocols, and is the main reason why polynomial backoff is stable for any arrival rate $\lambda < 1$.

Let us consider next why the potential is expected to decrease in Case 2. A good property of the potential function in Equation (3.1) (which is satisfied by polynomial backoff but not exponential backoff as we will see later) is that whenever we have a counter (say b_i) which is very large (i.e. larger than B) then the expected decrease in its contribution to the potential function in a

single step is

$$(b_i + 1)^{-\alpha}(b_i + 1)^{\alpha+1/2} = (b_i + 1)^{1/2} \geq \sqrt{B + 1}.$$

This can result in a significant drop in the potential function since B can be set to be very large.

On the other hand, even in a blocked channel where the counters just keep growing and no user succeeds, the expected change in the contribution of counter b_i to the potential is

$$\begin{aligned} & E[(b_i(t+1) + 1)^{\alpha+1/2} \mid \text{Channel is blocked}] - (b_i(t) + 1)^{\alpha+1/2} \\ &= (b_i(t) + 2)^{\alpha+1/2}(b_i(t) + 1)^{-\alpha} + (b_i(t) + 1)^{\alpha+1/2}(1 - (b_i(t) + 1)^{-\alpha}) \\ &\quad - (b_i(t) + 1)^{\alpha+1/2} \\ &= (b_i(t) + 2)^{\alpha+1/2}(b_i(t) + 1)^{-\alpha} - (b_i(t) + 1)^{1/2} \\ &= (b_i(t) + 1)^{1/2} \left((1 + 1/(b_i(t) + 1))^{\alpha+1/2} - 1 \right) \\ &= (b_i(t) + 1)^{-1/2} \left(\frac{(1 + 1/(b_i(t) + 1))^{\alpha+1/2} - 1}{1/(b_i(t) + 1)} \right) \\ &\leq \frac{(\alpha + 1/2)(1 + 1/(b_i(t) + 1))^{\alpha-1/2}}{\sqrt{b_i(t) + 1}} \\ &\leq \frac{(\alpha + 1/2) \cdot 2^{\alpha-1/2}}{\sqrt{b_i(t) + 1}}. \end{aligned}$$

The first equation is explained as follows. In a blocked channel where no user succeeds, user i sends (and its backoff counter will increase by 1) with probability $(b_i(t) + 1)^{-\alpha}$, and stays silent (and the backoff counter stays the same) with probability $1 - (b_i(t) + 1)^{-\alpha}$. The first inequality is true by the following inequality which can be shown using the Mean Value Theorem. For any h and constant $\beta > 1$,

$$\frac{(1+h)^\beta - 1}{h} \leq \beta(1+h)^{\beta-1}.$$

Thus the expected increase in the backoff counter part of the potential is bounded from above by $\Theta(n)$, while the expected decrease is $\sqrt{B}$. By choosing B to be large enough, we can easily show that the expected change in the potential is negative [27].

Although capture is very visible in simulations for binary and c -ary exponential backoff, this type of potential function and case division *does not* work for exponential backoff. We argue below that given the same cases, there is no potential function similar to the one used by Håstad, Leighton, and Rogoff [27] that can drop in Case 2 for exponential backoff.

In particular, let us define the potential function for c -ary exponential backoff as

$$f(X_t) = c_1 \sum_{i=1}^n q_i + \sum_{i=1}^n c^{b_i} g(b_i).$$

The argument in Case 2 depends on the fact that the expected decrease for the user with the big backoff counter is

$$c^{-b_i} \times c^{b_i} g(b_i) = g(b_i)$$

for some monotone increasing function $g(b_i)$. The function $g(b_i)$ naturally may be monotone increasing since we will increase $b_i = B$ to a value which will give us a suitable drop in the potential.

At the same time, the expected increase in the contribution of the rest of the backoff counters must not exceed a constant c' . In particular, the expected increase for user i must be

$$\begin{aligned} c^{-b_i} c^{b_i+1} g(b_i + 1) + (1 - c^{-b_i}) c^{b_i} g(b_i) - c^{b_i} g(b_i) &= \\ c g(b_i + 1) - g(b_i) &\leq c'. \end{aligned}$$

The last inequality cannot be satisfied by any monotone increasing function $g(\cdot)$ since $c > 1$. Therefore, we cannot use a potential function and case division similar the one used by Håstad et al. [27]. In the next section, we will see that a potential function with $g(b_i) = 1$ and a different case partition will enable us to show that the drift is negative if we run for an adequate number of steps.

3.3 The Stability of c -ary Exponential Backoff

In this section we prove the main result of this chapter regarding the stability of c -ary exponential backoff for the finite model. As a special case, when $c = 2$,

this result improves the best previous bound on the stability rate for the binary exponential backoff protocol [25], which was $\lambda < n^{-\Theta(\log n)}$, by increasing it to $\lambda \leq \Theta(n^{-(.75+\epsilon)})$ for any $\epsilon > 0$.

Theorem 3.1 *For any $c > 1$, there exists a positive constant α such that the c -ary exponential backoff protocol running on a system of n users is positive recurrent for symmetric arrivals $\lambda_i = \lambda/n$, sufficiently large n , and an arrival rate $\lambda \leq \frac{1}{\alpha n^{1-\eta}}$, for any $\eta < 0.25$.*

The proof of this theorem is presented below. Positive recurrence will be shown using the multiple step version of Foster's theorem (Theorem 2.2) which is due to Fayolle, Malyshev, and Menshikov. The finite set C is defined to be the set containing the empty state only (i.e. the state with no messages in the queues and with all counters initialised to zero).

Let us start by stating the potential function that will be used to prove this theorem. Define μ to be a constant in the range $(\eta, 0.5 - \eta)$. For any state X_t , let

$$f(X_t) = \alpha c n^{2-\eta-\mu} \sum_{i=1}^n q_i(t) + \sum_{i=1}^n c^{b_i(t)}. \quad (3.2)$$

The function k in Theorem 2.2 is defined as $k((0, 0, \dots, 0)) = 1$ for the empty state. For the states $s \in \mathbf{X} - C$, the function k will be defined below according to the starting state. Equation (2.5) is clearly satisfied by our choice of C and k . To see this, note that for $s \in C$ (i.e. s is the empty state),

$$\begin{aligned} E[f(X_{t+k(X_t)}) \mid X_t = s] &\leq \alpha c n^{2-\mu-\eta} \lambda + n \\ &< \infty. \end{aligned}$$

To show positive recurrence for the Markov chain X_t , we must show that Inequality (2.4) is satisfied. The proof of this will be divided into three cases according to the number of active users with “small” backoff counters. The number of steps that the system will run before we get the drop in the potential will depend on the case it starts from. However, this will be bounded by a constant

$$K = 2c^2 n \log_c(n).$$

To present the cases we need to define some more notation. Let $\beta = \frac{c}{c-1} + 1$, and for any state X_t , let

- $m(X_t)$ = the number of users i with $q_i(t) > 0$ and $b_i(t) < \log_c \beta + \log_c n$, and
- $m'(X_t)$ = the number of users i with $q_i(t) > 0$ and $b_i(t) < (1 - \mu - \eta) \log_c n + 1$.

We need only consider the states in $\mathbf{X} - C$, which will be divided into the following cases:

- Case 1. $m'(X_t) = 0$ and $m(X_t) < n^{1-\eta-\mu}$.
- Case 2. $m(X_t) \geq n^{1-\eta-\mu}$ or $m'(X_t) > n^{0.4}$.
- Case 3. $0 < m'(X_t) \leq n^{0.4}$ and $m(X_t) < n^{1-\eta-\mu}$.

For each case the value of the function k will be specified, and we will show that starting from state X_t ,

$$E[f(X_{t+k(X_t)}) - f(X_t)] \leq -\epsilon k(X_t).$$

Note that we will set $\epsilon = 1 - 2/\alpha$.

Case 1: $m'(X_t) = 0$ and $m(X_t) < n^{1-\eta-\mu}$.

For every state s such that $m'(s) = 0$ and $m(s) < n^{1-\eta-\mu}$ we define $k(s) = 1$. We wish to show that, if $s \in \mathbf{X} - C$ and $X_t = s$, then

$$E[f(X_{t+1}) - f(X_t)] \leq -\epsilon.$$

First, we give some intuition as to why the potential f is expected to drop in a single step. In this case (since $m'(X_t) = 0$) all users which have messages to send have *large* backoff counters. Furthermore (since $m(X_t) < n^{1-\eta-\mu}$) most backoff counters (all but at most $n^{1-\eta-\mu}$) are very large. This means that collisions are fairly unlikely. The expected drop in f mainly comes from the fact that if user i does send (which happens with probability c^{-b_i}) and succeeds (which is fairly likely since other users are likely to stay silent), then f drops by $c^{b_i} - 1$. The proof that f is expected to go down comes from a careful analysis of a single step and uses the same general approach as the one used in the proof of Lemma 5.7 in [27].

For convenience, we use m as shorthand for $m(X_t)$ and we use ℓ to denote the number of users i with $q_i(t) > 0$. Without loss of generality, we assume that these are users $1, \dots, \ell$. We use p_i to denote the probability that user i sends on step t . So $p_i = c^{-b_i(t)}$ if $i \in [1, \dots, \ell]$ and $p_i = \lambda/n$ otherwise.

The probability that all users stay silent at the current step is denoted by

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

We will also need to define the sum

$$S = \sum_{i=1}^n \frac{p_i}{1 - p_i}.$$

Note that the expected number of successes at step t is $S \times T$.

Let $I_{a,i,t}$ be the 0/1 indicator random variable which is 1 iff there is an arrival at user i during step t and let $I_{s,i,t}$ be the 0/1 indicator random variable which is 1 iff user i succeeds in sending a message at step t .

Define σ_i as the probability that user i collides at step t and ρ_i as the prob-

ability that user i sends successfully at step t . Then

$$\begin{aligned}
& E[f(X_{t+1}) - f(X_t)] \\
&= \alpha c n^{2-\eta-\mu} \sum_{i=1}^n (E[I_{a,i,t}] - E[I_{s,i,t}]) + \sum_{i=1}^n (E[c^{b_i(t+1)}] - c^{b_i(t)}), \\
&= \alpha c n^{2-\eta-\mu} (\lambda - ST) + \sum_{i=1}^n ((c-1)c^{b_i(t)}\sigma_i - (c^{b_i(t)} - 1)\rho_i), \quad (3.3) \\
&= \alpha c n^{2-\eta-\mu} (\lambda - ST) + \\
&\quad \sum_{i=1}^n \left((c-1)c^{b_i(t)}p_i \left(1 - \frac{T}{1-p_i}\right) - (c^{b_i(t)} - 1)p_i \frac{T}{1-p_i} \right), \\
&= \alpha c n^{2-\eta-\mu} (\lambda - ST) \\
&\quad + \sum_{i=1}^{\ell} (c-1) \left(1 - \frac{T}{1-p_i}\right) + \sum_{i=\ell+1}^n (c-1) \frac{\lambda}{n} \left(1 - \frac{T}{1-p_i}\right) - \ell T, \\
&= \alpha c n^{2-\eta-\mu} (\lambda - ST) + \left(\ell + \frac{(n-\ell)\lambda}{n}\right)(c-1) \\
&\quad - T(c-1) \left(\sum_{i=1}^{\ell} \frac{1}{1-p_i} + \sum_{i=\ell+1}^n \frac{p_i}{1-p_i} \right) - \ell T, \\
&= \alpha c n^{2-\eta-\mu} (\lambda - ST) + \left(\ell + \frac{(n-\ell)\lambda}{n} - ST - \ell T\right)(c-1) - \ell T, \\
&= \alpha c n^{2-\eta-\mu} \lambda + (c-1)\ell + (c-1) \frac{(n-\ell)\lambda}{n} \\
&\quad - T((\alpha c n^{2-\eta-\mu} + c-1)S + c\ell), \quad (3.4)
\end{aligned}$$

To see why Equality (3.3) holds, note that with probability σ_i , $b_i(t+1) = b_i(t) + 1$, with probability ρ_i , $b_i(t+1) = 0$, and otherwise, $b_i(t+1) = b_i(t)$. So, when a user sends and collides, the change in the potential is $c^{b_i(t)+1} - c^{b_i(t)}$, and if it succeeds then the change is $-(c^{b_i(t)} - c^0)$ because the backoff counter is reset to zero. We now find lower bounds for S and T . First,

$$\begin{aligned}
S &= \sum_{i=1}^n \frac{p_i}{1-p_i} \\
&= \sum_{i=1}^{\ell} \left(\frac{c^{-b_i(t)}}{1 - c^{-b_i(t)}} \right) + \frac{\lambda(n-\ell)}{n-\lambda} \\
&\geq \sum_{i=1}^m \left(\frac{1}{\beta n - 1} \right) + \frac{\lambda(n-\ell)}{n-\lambda} \\
&= \frac{m}{\beta n - 1} + \frac{\lambda(n-\ell)}{n-\lambda}. \quad (3.5)
\end{aligned}$$

Next,

$$\begin{aligned}
T &= \prod_{i=1}^n (1 - p_i) \\
&\geq \left(1 - \frac{1}{cn^{1-\eta-\mu}}\right)^m \left(1 - \frac{1}{\beta n}\right)^{\ell-m} \left(1 - \frac{\lambda}{n}\right)^{n-\ell} \\
&\geq 1 - \frac{m}{cn^{1-\eta-\mu}} - \frac{\ell-m}{\beta n} - \frac{\lambda(n-\ell)}{n}. \tag{3.6}
\end{aligned}$$

Combining Equations (3.4), (3.5) and (3.6), we get the following upper bound.

$$\begin{aligned}
E[f(X_{t+1}) - f(X_t)] &\leq \\
&\alpha c n^{2-\eta-\mu} \lambda + \left(\ell + \frac{(n-\ell)\lambda}{n}\right)(c-1) - \left(1 - \frac{m}{cn^{1-\eta-\mu}} - \frac{\ell-m}{\beta n} - \frac{\lambda(n-\ell)}{n}\right) \times \\
&\quad \left((\alpha c n^{2-\eta-\mu} + c - 1) \left(\frac{m}{\beta n - 1} + \frac{\lambda(n-\ell)}{n-\lambda}\right) + c \ell\right). \tag{3.7}
\end{aligned}$$

For the sake of the analysis below, assume that the arrival rate is $\lambda = 1/(\alpha' n^{1-\eta})$ for some $\alpha' \geq \alpha$. We will let $g(m, \ell)$ be the quantity in Equation (3.7) plus ϵ and we will show that $g(m, \ell)$ is negative for all values of $0 \leq m < n^{1-\eta-\mu}$ and all $\ell \geq m$. In particular, for every fixed positive value of m , we will show that

1. $g(m, m)$ is negative,
2. $g(m, n)$ is negative, and
3. $\frac{\partial^2}{\partial \ell^2} g(m, \ell) > 0$. ($g(m, \ell)$ is concave up as a function of ℓ for the fixed value of m so $g(m, \ell)$ is negative for all $\ell \in [m, n]$.)

We will handle the case $m = 0$ similarly except that $m = \ell = 0$ corresponds to the start state, which is excluded because it belongs to the set C . Therefore, we will replace Item 1 with the following for $m = 0$.

- 1'. $g(0, 1)$ is negative.

The details of the proof are now merely calculations and are presented as follows.

1. $g(m, m)$ is negative:

$$\begin{aligned}
& g(m, m) \times c \alpha' n^{(2-2\eta)} (\beta n - 1) (\alpha' n^2 - n^\eta) = \\
& c m - c^2 m - c^2 m^2 - c n + c^2 n - c m n + 2 c^2 m n + c n^2 \\
& - c^2 n^2 + c^2 n^{3-\eta-\mu} \alpha + c^2 m n^{3-\eta-\mu} \alpha - c^2 n^{4-\eta-\mu} \alpha - c m n \beta \\
& + c^2 m n \beta - c m^2 n \beta + 2 c^2 m^2 n \beta + c n^2 \beta - c^2 n^2 \beta + 2 c m n^2 \beta \\
& - 3 c^2 m n^2 \beta - c n^3 \beta + c^2 n^3 \beta + c^2 m^2 n^{3-\eta-\mu} \alpha \beta - c^2 n^{4-\eta-\mu} \alpha \beta \\
& - 2 c^2 m n^{4-\eta-\mu} \alpha \beta + c^2 n^{5-\eta-\mu} \alpha \beta - 2 c m n^{2-\eta} \alpha' + c^2 m n^{2-\eta} \alpha' \\
& + c m^2 n^{2-\eta} \alpha' - c m n^{3-\eta} \alpha' + c m^2 n^{1+\mu} \alpha' + m n^{2+\mu} \alpha' - c m n^{2+\mu} \alpha' \\
& - c m n^{4-\eta} \alpha \alpha' - c^2 m^2 n^{4-2\eta-\mu} \alpha \alpha' + c^2 m n^{5-2\eta-\mu} \alpha \alpha' \\
& + c m n^{3-\eta} \beta \alpha' - c^2 m^2 n^{3-\eta} \beta \alpha' + c^2 m n^{4-\eta} \beta \alpha' + m^2 n^{2+\mu} \beta \alpha' \\
& - 2 c m^2 n^{2+\mu} \beta \alpha' - m n^{3+\mu} \beta \alpha' + c m n^{3+\mu} \beta \alpha' - c m^2 n^{4-\eta} \alpha \beta \alpha' \\
& + c m n^{5-\eta} \alpha \beta \alpha' + c^2 m n^{5-2\eta-\mu} \alpha \beta \alpha' + c n^{2-\eta} \epsilon \alpha' - c n^{3-\eta} \beta \epsilon \alpha' \\
& + 2 c m n^{4-2\eta} (\alpha')^2 - c^2 m n^{4-2\eta} (\alpha')^2 - m^2 n^{3-\eta+\mu} (\alpha')^2 \\
& + c m^2 n^{5-2\eta} \alpha (\alpha')^2 - c^2 m n^{6-3\eta-\mu} \alpha (\alpha')^2 - c m n^{5-2\eta} \beta (\alpha')^2 \\
& + c m^2 n^{4-\eta+\mu} \beta (\alpha')^2 - c n^{4-2\eta} \epsilon (\alpha')^2 + c n^{5-2\eta} \beta \epsilon (\alpha')^2
\end{aligned}$$

Since $\eta + \mu < .5$, the dominant term is $-c^2 n^{(6-3\eta-\mu)} \alpha m \alpha'^2$. Note that there is a positive term (namely, $c m^2 n^{(5-2\eta)} \alpha (\alpha')^2$) which could be $1/c$ as big as the negative term if m is as big as $n^{1-\eta-\mu}$ (the upper bound for Case 1), but all other terms are asymptotically smaller.

2. $g(m, n)$ is negative:

$$\begin{aligned}
g(m, n) \times c \alpha' \beta n (\beta n - 1) = & \\
& -c^2 n^{2-\mu} \alpha \beta + c^2 n^{3-\mu} \alpha \beta^2 + c m^2 \alpha' - c^2 m^2 \alpha' - c m n \alpha' \\
& + 2 c^2 m n \alpha' - c^2 n^2 \alpha' - c^2 m^2 n^{2-\eta-\mu} \alpha \alpha' + c^2 m n^{3-\eta-\mu} \alpha \alpha' \\
& + c m n \beta \alpha' - c^2 m n \beta \alpha' + c n^2 \beta \alpha' - c^2 m n^2 \beta \alpha' + c^2 n^3 \beta \alpha' \\
& - m^2 n^{\eta+\mu} \beta \alpha' + c m^2 n^{\eta+\mu} \beta \alpha' - c m n^{1+\eta+\mu} \beta \alpha' + c m^2 n^2 \alpha \beta \alpha' \\
& - c^2 m n^{3-\eta-\mu} \alpha \beta \alpha' - c n^3 \beta^2 \alpha' + c m n^{2+\eta+\mu} \beta^2 \alpha' - c n \beta \epsilon \alpha' \\
& + c n^2 \beta^2 \epsilon \alpha'
\end{aligned}$$

Since $\eta + \mu < .5$ and $\beta \geq \frac{c}{c-1} + 1$, the term $-cn^3\beta^2\alpha'$ dominates $+c^2n^3\beta\alpha'$. For the same reason, the term $-c^2 mn^{(3-\eta-\mu)}\alpha\alpha'\beta$ dominates the two positive terms $+c^2 mn^{(3-\eta-\mu)}\alpha\alpha'$ and $+c m^2n^2\alpha\beta\alpha'$. The other terms are all asymptotically smaller.

3. $\frac{\partial^2}{\partial \ell^2} g(m, \ell) > 0$:

$$\frac{\partial^2}{\partial \ell^2} g(m, \ell) = 2 \left(\frac{1}{n \beta} - \frac{1}{\alpha' n^{2-\eta}} \right) \left(\frac{cn^2\alpha' + n^\eta - 2c n^\eta - cn^{2-\eta}\alpha'}{\alpha' n^2 - n^\eta} \right)$$

1'. $g(0, 1)$ is negative:

$$\begin{aligned}
g(0, 1) \times \alpha' \beta n^{(2-2\eta)} (\alpha' n^2 - n^\eta) = & \\
& -2 \beta + 3 c \beta + 3 n \beta - 4 c n \beta - n^2 \beta + c n^2 \beta + c n^{2-\eta-\mu} \alpha \beta \\
& - 3 c n^{3-\eta-\mu} \alpha \beta + c n^{4-\eta-\mu} \alpha \beta + n^{1-\eta} \alpha' - 2 c n^{1-\eta} \alpha' - n^{2-\eta} \alpha' \\
& + c n^{2-\eta} \alpha' - c n^{3-2\eta-\mu} \alpha \alpha' + c n^{4-2\eta-\mu} \alpha \alpha' + n^{2-\eta} \beta \alpha' - c n^{2-\eta} \beta \alpha' \\
& + c n^{3-\eta} \beta \alpha' + c n^{4-2\eta-\mu} \alpha \beta \alpha' - n^{2-\eta} \beta \epsilon \alpha' + c n^{3-2\eta} (\alpha')^2 \\
& - n^{4-2\eta} \beta (\alpha')^2 + n^{4-2\eta} \beta \epsilon (\alpha')^2
\end{aligned}$$

Since $\alpha'(1-\epsilon) \geq \alpha(1-\epsilon) > 1$, $\eta + \mu < .5$, and $\mu > \eta$, the term $-n^{(4-2\eta)}\beta(\alpha')^2(1-\epsilon)$ dominates the term $+cn^{(4-\eta-\mu)}\alpha\beta$. The other terms are asymptotically smaller.

Case 2: $m(X_t) \geq n^{1-\eta-\mu}$ or $m'(X_t) > n^4$.

For every state s such that $m(s) \geq n^{1-\eta-\mu}$ or $m'(s) > n^4$, we will define an integer k (which depends upon the starting state s) and we will show that, if $X_t = s$, then

$$E[f(X_{t+k}) - f(X_t)] \leq -\epsilon k,$$

where $\epsilon = 1 - 2/\alpha$. This will be more than enough for satisfying Equation (2.4).

For convenience, we will use m as shorthand for $m(X_t)$ and m' as shorthand for $m'(X_t)$. If $m \geq n^{1-\eta-\mu}$ then we will define $r = m$, $W = n^{\eta+\mu} \lceil \log_c r \rceil c^{-8}$, $A = W$, $b = \log_c \beta + \log_c n$ and $v = n$. Otherwise, we will define $r = m'$, $W = \lceil \log_c r \rceil c^{-8}$, $A = 0$, $b = (1 - \eta - \mu) \log_c n + 1$, and $v = c \lceil n^{1-\eta-\mu} \rceil$. In either case, we will define $k = c^2(r + v) \lceil \log_c r \rceil$.

The intuition behind the proof is as follows. First, since many users have small backoff counters, it is fairly likely that a collision occurs on the first step. So we do not expect the potential f to drop in a single step. Instead, we study the evolution of the system over k steps. With sufficiently high probability, the backoff counters get driven up during the first $\Theta(r \log r)$ steps. We refer to these steps as “the preamble”. During the remaining steps, the backoff counters stay reasonably high except during steps which occur shortly after

1. arrivals (but there are likely to be few of these since we only run for k steps), and
2. successful sends (which help to reduce f).

We refer to these as “exceptional steps”. Without loss of generality, there are few of them, since otherwise there are many successes and the potential goes down. Although the backoff counters stay high (as we just explained), most of them do not get too high, since we only run for k steps. So the probability of success during any given step which is not exceptional or in the preamble is high. Finally, with sufficiently high probability, there are at least W successes, and this reduces the potential.

A technical difficulty in the proof is clarifying the independence between

some of the events and for this reason, it is helpful to identify “preamble steps” (steps in τ_0), “exceptional steps” (steps in τ_1), and also “following steps”. The formal definition of “following steps” will be given later. Typically, these steps follow at least W successes. Informally, we will stop counting successes in the following steps because we already counted the W successes needed to show that the potential drops.

Let τ be the set of all steps $\{t, \dots, t + k - 1\}$ and let S be the random variable which denotes the number of successes that the system has during τ . Let p denote $\Pr(S \geq W)$. Then we have

$$\begin{aligned} E[f(X_{t+k}) - f(X_t)] &\leq \alpha c n^{2-\eta-\mu} (\lambda k - E[S]) + \sum_{i=1}^n \sum_{t'=t+1}^{t+k} E[c^{b_i(t')} - c^{b_i(t'-1)}] \\ &\leq \alpha c n^{2-\eta-\mu} (\lambda k - Wp) + kn(c - 1) \\ &\leq -\epsilon k, \end{aligned}$$

where the final inequality holds as long as $\alpha p \geq 2^{10} c^{13}$ and n is sufficiently big. Note also that $E[c^{b_i(t')} - c^{b_i(t'-1)}] \leq c - 1$ since in an externally-jammed channel (where no user succeeds), the expected change in c^{b_i} is $c - 1$.

To show that the last inequality is true, consider the following two cases which are divided according to the values of m and m' . In both cases, the constant α is chosen so that

$$\alpha p \geq 8c^{12}.$$

This allows the expected number of successful messages to dominate the expected change in the potential, and thus makes it negative.

Subcase A: $m \geq n^{1-\eta-\mu}$

In this case, $r = m$, $W = n^{\eta+\mu} \lceil \log_c r \rceil c^{-8}$, $v = n$, and $k = c^2(r + v) \lceil \log_c r \rceil$. Since

$k \leq 2c^2 n \lceil \log_c m \rceil$, the expected change in the potential in k steps is at most

$$\begin{aligned}
& \alpha c n^{2-\eta-\mu} \lambda k - \alpha c n^{2-\eta-\mu} W p + k n (c - 1) \\
& \leq \alpha c n^{2-\eta-\mu} (\alpha' n^{1-\eta})^{-1} k - \alpha c n^{2-\eta-\mu} W p + k n (c - 1) \\
& \leq \frac{\alpha}{\alpha'} n^{1-\mu} k - \alpha c n^{2-\eta-\mu} n^{\eta+\mu} \lceil \log_c r \rceil c^{-8} p + k n c \\
& \leq 2c^2 n^{2-\eta} \lceil \log_c m \rceil - \alpha c^{-7} n^2 \lceil \log_c m \rceil p + 2c^3 n^2 \lceil \log_c m \rceil \\
& \leq 2c^2 n^{2-\mu} \lceil \log_c m \rceil - 8c^3 n^2 \lceil \log_c m \rceil + 2c^3 n^2 \lceil \log_c m \rceil \\
& \leq -4c^3 n^2 \lceil \log_c m \rceil \\
& \leq -2c n k \\
& \leq -\epsilon k
\end{aligned}$$

when n is sufficiently large.

Subcase B: $m < n^{1-\eta-\mu}$, $m' > n^4$

In this case, $r = m'$, $W = \lceil \log_c r \rceil c^{-8}$, and $v = c \lceil n^{1-\eta-\mu} \rceil$, and $k = c^2(r+v) \lceil \log_c r \rceil$. Note that by definition, $m' < n^{1-\eta-\mu}$. Then, since $k \leq 2c^3 \lceil n^{1-\eta-\mu} \rceil \lceil \log_c m' \rceil$, for large n ,

$$\begin{aligned}
& \alpha c n^{2-\eta-\mu} \lambda k - \alpha c n^{2-\eta-\mu} W p + k n (c - 1) \\
& \leq \alpha c n^{2-\eta-\mu} (\alpha' n^{1-\eta})^{-1} k - \alpha p c n^{2-\eta-\mu} (\log_c m' c^{-8}) + k n c \\
& \leq c n^{1-\mu} 2c^3 \lceil n^{1-\eta-\mu} \rceil \lceil \log_c m' \rceil - \alpha p c n^{2-\eta-\mu} \lceil \log_c m' \rceil c^{-8} + 2c^3 n^{2-\eta-\mu} \lceil \log_c m' \rceil c \\
& \leq 4c^3 c \lceil n^{2-\eta-2\mu} \rceil \lceil \log_c m' \rceil - 8c^4 n^{2-\eta-\mu} \lceil \log_c m' \rceil + 2c^3 c n^{2-\eta-\mu} \lceil \log_c m' \rceil \\
& \leq -2c^4 n^{2-\eta-\mu} \lceil \log_c m' \rceil \\
& \leq -n k \\
& \leq -\epsilon k.
\end{aligned}$$

Thus, it suffices to find a positive lower bound for p which is independent of n . We do this with plenty to spare. In particular, we show that $p \geq 1-5 \times 10^{-5}$.

We start with a technical lemma, which describes the behaviour of a single user with high probability. This lemma is needed because the exact values of

the backoff counters of the users after the start state are unknown. However, we can estimate the value of a counter with high probability using the following lemma.

Lemma 3.2 *Let j be a positive integer, and let δ be a positive integer which is at least c . Suppose that $q_i(t) > 0$. Then, with probability at least $1 - \frac{\lceil \log_c j \rceil}{j^{\delta/(c \ln c)}}$, either user i succeeds in at least one of the steps in the interval $[t, \dots, t + \delta j \lceil \log_c j \rceil - 1]$, or $b_i(t + \delta j \lceil \log_c j \rceil) \geq \lceil \log_c j \rceil$.*

Proof: Suppose that user i is running in an externally-jammed channel (so every send results in a collision). Let X_z denote the number of steps $t' \in [t, \dots, t + \lceil \delta j \log_c(j) \rceil]$ with $b_i(t') = z$. We claim that

$$\Pr(X_z > \delta \lceil \log_c j \rceil c^{z-1}) < j^{-\delta/(c \ln c)}.$$

This proves the lemma since $\sum_{z=0}^{\lceil \log_c j \rceil - 1} \delta \lceil \log_c j \rceil c^{z-1} \leq \delta j \lceil \log_c j \rceil$.

To prove the claim, note that $X_0 \leq 1$, so $\Pr(X_0 > \delta \lceil \log_c j \rceil c^{-1}) = 0 < j^{-\delta/(c \ln c)}$. For $z > 0$, note that

$$\Pr(X_z > \delta \lceil \log_c j \rceil c^{z-1}) \leq (1 - c^{-z})^{\delta \lceil \log_c j \rceil c^{z-1}} < j^{-\delta/(c \ln c)}.$$

□

Next, we define some events. We will show that the events are likely to occur, and, if they do occur, then $\mathcal{S}$ is likely to be at least W . This will allow us to conclude that $p \geq 1 - 5 \times 10^{-5}$, which will finish Case 2. We start by defining $B = \lceil W \rceil + \lceil A \rceil$, $k' = c^2 r \lceil \log_c r \rceil$, and $k'' = c^2 B \lceil \log_c B \rceil$. Next, we give names to some of the steps in $\tau = \{t, \dots, t + k - 1\}$. Let τ_0 be the preamble of τ consisting of steps $\{t, \dots, t + k' - 1\}$. For every i , let $\tau'(i)$ be the set of times in τ when user i will “definitely” send (by definitely we mean that the users will send with probability 1.) In particular, $t' \in \tau'(i)$ if and only if

1. $b_i(t') = 0$ and $q_i(t') > 0$, or
2. $b_i(t') = 0$ and there is an arrival at user i at t' .

Proof: The expected number of arrivals in τ is λk . If $m \geq n^{1-\eta-\mu}$, then $A = n^{\eta+\mu} \lceil \log_c r \rceil c^{-8} \geq 2\lambda k$. By a Chernoff bound, the probability that there are this many arrivals is at most $e^{-\lambda k/3} \leq 10^{-5}$. Otherwise, $A = 0$ and $\lambda k = o(1)$. Thus, $\Pr(\mathbf{E1}) \geq (1 - \lambda/n)^{nk} \geq 1 - \lambda k \geq 1 - 10^{-5}$. $\square$

Lemma 3.4 *If n is sufficiently large then $\Pr(\overline{\mathbf{E2}}) \leq 10^{-5}$.*

Proof: Apply Lemma 3.2 to each of the r users with $\delta = c^2$ and $j = r$. Then $\Pr(\overline{\mathbf{E2}}) \leq r \frac{\lceil \log_c r \rceil}{r c / (\ln c)} \leq 10^{-5}$. $\square$

Lemma 3.5 *If n is sufficiently large then $\Pr(\overline{\mathbf{E3}}) \leq 10^{-5}$.*

Proof: Note that $k \leq 4c^2 v \log_c r$. Also note that the probability of a given user i sending at step t' when $b_i(t') \geq b + \lceil \log_c \log_c(r) \rceil + 1 + \log_c 16$ is at most $1/(16c^2 v \log_c r)$. Thus the probability that user i sends at all in the k steps of τ is at most $1/4$. By a Chernoff bound, the probability that over half of the r users with $q_i(t) > 0$ and $b_i(t) < b$ send when $b_i(t') \geq b + \lceil \log_c \log_c(r) \rceil + 1 + \log_c 16$ for some $t' \in \tau$ is at most $e^{-r/12} < 10^{-5}$. $\square$

Lemma 3.6 *If n is sufficiently large then $\Pr(\overline{\mathbf{E4}}) \leq 10^{-5}$.*

Proof: We can apply Lemma 3.2 separately to each of the (up to B) pairs (t', i) with $\delta = c^2$ and $j = B$. The probability that event $\mathbf{E4}$ does not hold is at most $\frac{B \lceil \log_c B \rceil}{B c / (\ln c)} \leq 10^{-5}$. $\square$

We now wish to show that $\Pr(\mathcal{S} < W \mid \mathbf{E1} \wedge \mathbf{E2} \wedge \mathbf{E3} \wedge \mathbf{E4}) \leq 10^{-5}$. We begin with the following lemma.

Lemma 3.7 *Given any fixed sequence of states $X_t, \dots, X_{t+z}$ which does not violate $\mathbf{E2}$ or $\mathbf{E4}$, and satisfies $t + z \in \tau - \tau_0 - \tau_1 - \tau_2$, $q_i(t + z) > 0$, and $b_i(t + z) \leq b + \lceil \log_c \log_c(r) \rceil + 2 + \log_c 16$, the probability that user i succeeds at step $t + z$ is at least $\frac{1}{c^3 2^{12} c^b \log_c r}$.*

Proof: The conditions in the lemma imply the following.

- There are no users j with $q_j(t+z) > 0$ and $b_j(t+z) < \lceil \log_c B \rceil$ (since E4 holds).
- There are at most B users j with $b_j(t+z) < \lceil \log_c r \rceil$ (since E2 holds and at most B users succeed or have new arrivals).
- There are at most $r+B$ users j with $b_j(t+z) < b$ (since r started that way and at most B succeed or have new arrivals).
- There are at most $\min\{m+B, n\}$ users j with $b_j(t+z) < \log_c \beta + \log_c n$ (for similar reasons).

Thus, the probability that user i succeeds is at least

$$\begin{aligned}
& c^{-(b+\lceil \log_c \log_c(r) \rceil + 2 + \log_c 16)} \left(1 - \frac{1}{B}\right)^B \left(1 - \frac{1}{r}\right)^r \left(1 - \frac{1}{c^b}\right)^{m-r} \left(1 - \frac{1}{\beta n}\right)^{n-\min\{m+B, n\}} \\
& \geq \frac{1}{c^b \log_c r \cdot c^3 16} \frac{1}{4} \frac{1}{4} \frac{1}{4} \left(1 - \frac{n - \min\{m+B, n\}}{\beta n}\right) \\
& \geq \frac{1}{c^3 2^{12} c^b \log_c r}.
\end{aligned}$$

□

Corollary 3.8 *Given any fixed sequence of states $X_t, \dots, X_{t+z}$ which does not violate E2, E3, or E4, and satisfies $t+z \in \tau - \tau_0 - \tau_1 - \tau_2$, and given that there were fewer than B successes in $t, \dots, t+z-1$, then the probability that some user succeeds at step $t+z$ is at least $\frac{(\tau/2)-B}{c^3 2^{12} c^b \log_c r} \geq \frac{r}{c^4 2^{15} v \log_c r}$.*

Proof: Since there were fewer than B successes, at least $r-B$ of the users i with $q_i(t) > 0$ and $b_i(t) < b$ have not succeeded before step $t+z$. Since E3 holds, at least $r/2 - B$ of these have $b_i(t+z) \leq b + \lceil \log_c \log_c(r) \rceil + 2 + \log_c 16$. For all i and i' , the event that user i succeeds at step $t+z$ is disjoint with the event that user i' succeeds at step $t+z$. Finally, note that $(r/2) - B > r/4$ and $c^b \leq 2c v$. □

Lemma 3.9 *If n is sufficiently large then $\Pr(S < W \mid E1 \wedge E2 \wedge E3 \wedge E4) \leq 10^{-5}$.*

Proof: If E1 is satisfied then τ_2 does not start until there have been at least W successes. Since $|\tau - \tau_0 - \tau_1| \geq k - k' - Bk'' \geq v \lceil \log_c r \rceil / 2$, Corollary 3.8 shows that the probability of having fewer than W successes is at most

the probability of having fewer than W successes in $v \lceil \log_c r \rceil / 2$ Bernoulli trials with success probability $\frac{r}{c^4 2^{15} v \log_c r}$. Since W is at most half of the expected number of successes, a Chernoff bound shows that the probability of having fewer than W successes is at most $\exp(-\frac{rv \lceil \log_c r \rceil}{c^4 2^{19} v \log_c r}) \leq 10^{-5}$. $\square$

We conclude Case 2 by observing that p is at least $1 - \Pr(\overline{E1}) - \Pr(\overline{E2}) - \Pr(\overline{E3}) - \Pr(\overline{E4}) - \Pr(\mathcal{S} < W \mid E1 \wedge E2 \wedge E3 \wedge E4)$. By Lemmas 3.3, 3.4, 3.5, 3.6, and 3.9, this is at least $1 - 5 \times 10^{-5}$.

Case 3: $0 < m'(X_t) \leq n^4$ and $m(X_t) < n^{1-\eta-\mu}$.

For every state s such that $0 < m'(s) \leq n^4$ and $m(s) < n^{1-\eta-\mu}$, we will define

$$k = \lceil 18 \frac{c \ln^2(c)}{\ln m'(s)} + 12c \ln c \rceil m'(s) \lceil \log_c m'(s) \rceil + \lceil n^{1-\eta-\mu} \rceil.$$

We will show that, if $X_t = s$, then $E[f(X_{t+k}) - f(X_t)] \leq -\epsilon k$.

The intuition behind the proof in this case is similar to that of Case 2 except that we do not have enough small backoff counters to achieve W successes (as in Case 2) even though we may have too many to make the potential drop in a single step (as in Case 1). We study the evolution of the system over k steps. The backoff counters are likely to be driven up in the first $\Theta(m' \log m')$ steps. After that, we are likely to have a single success, which is enough to make the potential drop.

Once again, we will use m as shorthand for $m(X_t)$ and m' as shorthand for $m'(X_t)$. Let $\tau = \{t, \dots, t+k-1\}$ and let $\mathcal{S}$ be the number of successes that the system has in τ . Let p denote $\Pr(\mathcal{S} \geq 1)$. As in Case 2,

$$E[f(X_{t+k}) - f(X_t)] \leq \alpha c n^{2-\eta-\mu} \lambda k - \alpha c n^{2-\eta-\mu} p + kn(c-1),$$

and this is at most $-\epsilon k$ as long as $\alpha p > \frac{c-1}{c}$. Thus, we will finish by finding a positive lower bound for p which is independent of n .

Since $m' > 0$, there is a user u such that $b_u(t) < (1 - \eta - \mu) \log_c n + 1$. Let $k' = 0$ if $m' = 1$, otherwise let

$$k' = \lceil 18 \frac{c \ln^2 c}{\ln m'} + 12c \ln c \rceil m' \lceil \log_c m' \rceil.$$

Let $\tau_0 = \{t, \dots, t + k' - 1\}$. We will now define some events, as in Case 2.

E1. There are no arrivals during τ .

E2. Every user i with $q_i(t) > 0$ and $b_i(t) < (1 - \eta - \mu) \log_c n + 1$ either sends successfully during τ_0 or has $b_i(t + k') \geq \lceil \log_c m' \rceil$.

E3. $b_u(t') < (1 - \eta - \mu) \log_c n + 12$ for all $t' \in \tau$.

Lemma 3.10 *If n is sufficiently large then $\Pr(\overline{E1}) \leq 2^{-10} c^{-18}$.*

Proof: As in the proof of Lemma 3.3, for sufficiently large n ,

$$\Pr(\text{E1}) \geq \left(1 - \frac{\lambda}{n}\right)^{nk} \geq 1 - \lambda k \geq 1 - \frac{2n^{1-\eta-\mu}}{\alpha n^{1-\eta}} \geq 1 - \frac{1}{2^{10} c^{18}}.$$

□

Lemma 3.11 $\Pr(\overline{E2}) \leq 2^{-10} c^{-18}$.

Proof: For $m' = 1$, $k' = 0$ by definition and $\Pr(\overline{E2}) = 0$. For $m' \geq 2$, we use Lemma 3.2 with $\delta = \lceil 18 \frac{c \ln^2(c)}{\ln m'} + 12c \ln c \rceil$ and $j = m'$, to get

$$\begin{aligned} \Pr(\overline{E2}) &\leq m' \cdot \frac{\lceil \log_c m' \rceil}{(m')^{\delta/(c \ln(c))}} \\ &\leq m' \cdot \frac{\lceil \log_c m' \rceil}{(m')^{18 \frac{\ln c}{\ln m'} + 12}} \\ &= m' \cdot \frac{\lceil \log_c m' \rceil}{(m')^{12}} \cdot \frac{1}{c^{18 \ln m' / \ln m'}} \\ &< 2^{-10} c^{-18}. \end{aligned}$$

□

Lemma 3.12 *If n is sufficiently large then $\Pr(\overline{E3}) \leq 2^{-10} c^{-18}$.*

Proof: For E3 to be violated, user u must make at least 12 attempts, one each with backoff counter $\lceil (1 - \eta - \mu) \log_c n + r \rceil$ for $r \in \{1, \dots, 12\}$. The probability

of this happening is

$$\begin{aligned}
\Pr(\overline{E3}) &\leq \binom{k}{12} \prod_{r=1}^{12} c^{-\lceil (1-\eta-\mu) \log_c n \rceil - r} \\
&\leq \left(\frac{ke}{12}\right)^{12} \left(\frac{1}{n^{1-\eta-\mu}}\right)^{12} c^{-\sum_{r=1}^{12} r} \\
&\leq \left(\frac{2en^{1-\eta-\mu}}{12n^{1-\eta-\mu}c^{6.5}}\right)^{12} \\
&\leq \left(\frac{1}{2c^{6.5}}\right)^{12} \\
&< 2^{-10} c^{-18}.
\end{aligned}$$

□

Lemma 3.13 *Given any fixed sequence of states $X_t, \dots, X_{t+z}$ which does not violate $E1$, $E2$, or $E3$ such that $t+z \in \tau - \tau_0$ and there are no successes during steps $[t, \dots, t+z-1]$, the probability that user u succeeds at step $t+z$ is at least $\frac{1}{2^5 c^{12} n^{1-\eta-\mu}}$.*

Proof: The conditions in the statement of the lemma imply the following.

- $q_u(t+z) > 0$ and $b_u(t+z) < (1-\eta-\mu) \log_c n + 12$.
- There are no users j with $b_j(t+z) < \lceil \log_c m' \rceil$.
- There are at most m' users j with $b_j(t+z) < (1-\eta-\mu) \log_c n + 1$.
- There are at most m users j with $b_j(t+z) < \log_c \beta + \log_c n$.
- There will be no arrivals on step $t+z$.

The probability of success for user u is at least

$$\begin{aligned}
&c^{-((1-\eta-\mu) \log_c n + 12)} \left(1 - \frac{1}{m'}\right)^{m'-1} \left(1 - \frac{1}{c n^{1-\eta-\mu}}\right)^{m-m'} \left(1 - \frac{1}{\beta n}\right)^{n-m} \\
&\geq \frac{1}{c^{12} n^{1-\eta-\mu}} \frac{1}{4} \frac{1}{4} \frac{1}{2} \\
&\geq \frac{1}{2^5 c^{12} n^{1-\eta-\mu}}.
\end{aligned}$$

□

Lemma 3.14 *If n is sufficiently large then $\Pr(S < 1 \mid E1 \wedge E2 \wedge E3) \leq 1 - 1/(2^6 c^{12})$.*

Proof: Lemma 3.13 implies that the probability of having no successes is at most the probability of having no successes in $|\tau - \tau_0|$ Bernoulli trials, each with success probability $\frac{1}{2^5 c^{12} n^{1-\eta-\mu}}$. Since $|\tau - \tau_0| \geq n^{1-\eta-\mu}$, this probability is at most

$$\left(1 - \frac{1}{2^5 c^{12} n^{1-\eta-\mu}}\right)^{n^{1-\eta-\mu}} \leq e^{-1/(2^5 c^{12})} \leq 1 - \frac{1}{2^6 c^{12}}.$$

□

We conclude Case 3 by observing that

$$p \geq 1 - \Pr(\overline{E1}) - \Pr(\overline{E2}) - \Pr(\overline{E3}) - \Pr(S < 1 \mid E1 \wedge E2 \wedge E3).$$

By Lemmas 3.10, 3.11, 3.12, and 3.14, this is at least

$$1 - 3 \times \frac{1}{2^{10} c^{18}} - \left(1 - \frac{1}{2^6 c^{12}}\right) \geq \frac{1}{2^7 c^{12}}.$$

Chapter 4

Instability Results for Exponential Backoff Protocols

We have shown in the previous chapter that the c -ary exponential backoff protocol (for any $c > 1$) is stable for some small enough arrival rate λ . In this chapter we will investigate the instability of the same protocol for $c \geq 2$, and we will show that all c -ary protocols are unstable for any $\lambda > \lambda_0 + 1/(4n - 2)$, where λ_0 is a constant approximately equal to 0.567. We will also examine binary exponential backoff for the special case when $n = 2$, and present some new results on its instability region.

Our definition of instability will be the following.

Definition 4.1 *A protocol is said to be unstable if $E[T_{ret}] = \infty$.*

At the level of the underlying Markov chain, this is equivalent to non-positivity. Thus, to show that a system is unstable, we will prove that the Markov chain is not positive recurrent. This in turn will prove that there is no stationary distribution for the Markov chain.

Note that this definition of instability is weaker than definitions requiring transience, where the probability of returning to the start state is < 1 . A transient chain is always non-positive, but a non-positive chain can be recurrent and thus non-transient [15]. Our definition of instability is also weaker than

the one adopted by Håstad, Leighton, and Rogoff [27] which requires showing that both $E[T_{ret}] = \infty$ and $E[L_{avg}] = \infty$.

To prove instability, we will again utilise a potential function and its drift. Basically, we must show that for any state with a large enough potential, the drift is non-negative. The jumps of the chain must also be bounded in expectation. This must be true for both negative and positive jumps. The following theorem will be used to prove instability.

Theorem 4.1 (Fayolle, Malyshev, Menshikov [15]) *An irreducible aperiodic time-homogeneous Markov chain X_t with state space $\mathbf{X}$ is not positive recurrent if there is a function $f : \mathbf{X} \rightarrow \mathbb{R}^+$, and constants V_0 and d such that*

1. *the sets $\{x \mid f(x) > V_0\}$ and $\{x \mid f(x) \leq V_0\}$ are nonempty, and*
2. *$E[f(X_{t+1}) - f(X_t) \mid X_t = x] \geq 0$ for all x with $f(x) > V_0$, and*
3. *$E[|f(X_{t+1}) - f(X_t)| \mid X_t = x] \leq d$ for every state x .*

4.1 Instability of c -ary Exponential Backoff

In this section we will obtain an instability result for c -ary exponential backoff for any $c \geq 2$. The model used in this section is the same as the one used for studying its stability region in Section 3.3. The arrival rates at the different users will be assumed to be symmetric and Bernoulli distributed.

For any constant $c \geq 2$, we will show that the c -ary exponential backoff protocol is unstable when the arrival rate is at most $\lambda_0 + \frac{1}{4n-2} \approx 0.567143 + \frac{1}{4n-2}$. The bound on λ and the proof method are the same as those introduced by Håstad, Leighton, and Rogoff [27] for studying the instability of binary exponential backoff. Even though the proof technique is not new, its application to the general version of exponential backoff is. It is hoped that this result will increase our understanding of the stability region of the general c -ary case. Certainly it shows that increasing the value of c will not produce protocols that are more stable than the binary case for such arrival rates.

We have seen a similar effect in the stability proof (in Section 3.3), where we showed the same stability result for any constant $c > 1$ (notice though that the stability result is asymptotic in n). It will be quite interesting to find out how and where the value of c affects the stability region of exponential backoff protocols.

Let λ_0 be the solution to $\lambda_0 = e^{-\lambda_0}$. The following is the main result of this section.

Theorem 4.2 *Suppose that c is a constant ≥ 2 . Then the c -ary exponential back-off protocol running on a system of n users with symmetric Bernoulli arrivals is unstable for any arrival rate $\lambda \geq \lambda_0 + \frac{1}{4n-2}$.*

Proof:

To investigate the instability of the c -ary exponential backoff protocol, we will use a potential function similar to the one used in Section 3.3. For any state X_t , define the potential function as

$$f(X_t) = (c-1)\alpha \sum_{i=1}^n q_i(t) + \sum_{i=1}^n c^{b_i(t)}, \quad (4.1)$$

where $\alpha = 2n-1$. The proof uses the same method used by Håstad, Leighton, and Rogoff [27] for showing that the binary exponential backoff is unstable for the same arrival rate.

Let ℓ denote the number of users i with $q_i(t) > 0$ (i.e. the number of nonempty users). Assume without loss of generality that these are the first ℓ users. For any user i , where $1 \leq i \leq \ell$, the probability that it will send is $p_i = c^{-b_i}$. For users $i > \ell$, which are empty at the start of the current step, the probability of sending is $p_i = \lambda_i$ (because new arrivals into empty queues are sent right away). Note that $\lambda_i = \lambda/n$ because we are assuming symmetric arrival rates.

To prove the theorem, we will use Theorem 4.1 to show that the underlying Markov chain is non-positive recurrent. By setting $V_0 = n$, Condition 1 of the theorem is satisfied since the set $\{x | f(x) \leq V_0\}$ contains the empty state.

Next, consider Condition 3 in Theorem 4.1 which is also true because of the following. We will examine the increase and decrease in the potential separately, and show that if they are multiplied by their probability (i.e. to calculate their expectations) then both are bounded by a constant d . Notice that the potential is divided into its queue part which is equal to $(c-1)\alpha \sum_{i=1}^n q_i(t)$ and its backoff counter part which is the term $\sum_{i=1}^n c^{b_i(t)}$. Note that the total potential is the sum of these two parts.

First consider the increase in the potential. The queue part of the potential can increase by at most $(c-1)\alpha \cdot n$ by having at most n arrivals into the queues in a single step. The backoff counter part is expected to increase by at most $n(c-1)$.¹ Thus the expected increase in the potential is bounded by $(c-1)\alpha \cdot n + n(c-1)$.

Next, consider the possible decrease in potential. We will calculate a bound on the expected decrease in the potential. The queue part of the potential can decrease in a single step by at most $(c-1)\alpha$, because only a single message can succeed in one step. When this happens, the expected decrease in the backoff part of the potential is at most $c^{-b_i} c^{b_i} = 1$. Thus, the expected decrease in the potential is at most $(c-1)\alpha + 1$. Hence,

$$E[|f(X_{t+1}) - f(X_t)| | X_t = x] \leq \max\{(c-1)\alpha \cdot n + n(c-1), (c-1)\alpha + 1\} \leq d.$$

This shows that Condition 3 is satisfied by the potential function f .

The rest of the proof is devoted to showing that the expected change in the potential is

$$E[f(X_{t+1}) - f(X_t)] \geq 0$$

for all states with $f(X_t) > V_0$. This shows that Condition 2 of Theorem 4.1 is satisfied, and hence the protocol is unstable. Following the method used in [27], this will be done by a detailed calculation of the expected change in a single step. The analysis will be divided into two cases.

¹This is obvious by examining the expected change in c^{b_i} for each i in a blocked channel. Assuming that no user succeeds, the expected change in c^{b_i} is $(c-1)$, and there are at most n such users.

Similar to what we did in Section 3.3, we will examine the expected change in the potential in a single step. Let T denote $\prod_{i=1}^n (1 - p_i)$ and let S denote $\sum_{i=1}^n \frac{p_i}{1-p_i}$. Note that the expected number of successes at step t is $S \times T$. Let $I_{a,i,t}$ be the 0/1 indicator random variable which is 1 iff there is an arrival at user i during step t and let $I_{s,i,t}$ be the 0/1 indicator random variable which is 1 iff user i succeeds in sending a message at step t .

Define σ_i as the probability that user i collides at step t and ρ_i as the probability that user i sends successfully at step t . Note that for $1 \leq i \leq \ell$,

$$\sigma_i = p_i \left(1 - \frac{T}{1 - p_i} \right) = c^{-b_i} \left(1 - \frac{T}{1 - c^{-b_i}} \right)$$

and for $\ell + 1 \leq i \leq n$,

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

Similarly, the probability that user i succeeds is $\rho_i = p_i \frac{T}{1-p_i}$.

Using these quantities, the expected change in a single step is

$$\begin{aligned}
& E[f(X_{t+1}) - f(X_t)] \\
&= (c-1)\alpha \sum_{i=1}^n (E[I_{a,i,t}] - E[I_{s,i,t}]) + \sum_{i=1}^n (E[c^{b_i(t+1)}] - c^{b_i(t)}), \\
&= (c-1)\alpha(\lambda - ST) + \sum_{i=1}^n ((c^{b_i(t)+1} - c^{b_i(t)})\sigma_i - (c^{b_i(t)} - 1)\rho_i), \\
&= (c-1)\alpha(\lambda - ST) + \\
&\quad \sum_{i=1}^n \left((c-1)c^{b_i(t)}p_i(1 - \frac{T}{1-p_i}) - (c^{b_i(t)} - 1)p_i\frac{T}{1-p_i} \right), \\
&= (c-1)\alpha(\lambda - ST) \\
&\quad + \sum_{i=1}^{\ell} (c-1)(1 - \frac{T}{1-p_i}) + \sum_{i=\ell+1}^n (c-1)\frac{\lambda}{n}(1 - \frac{T}{1-p_i}) - \ell T, \\
&= (c-1)\alpha(\lambda - ST) + (\ell + \frac{(n-\ell)\lambda}{n})(c-1) - \ell T \\
&\quad - T(c-1) \left(\sum_{i=1}^{\ell} \frac{1}{1-p_i} + \sum_{i=\ell+1}^n \frac{p_i}{1-p_i} \right), \\
&= (c-1) \left(\alpha(\lambda - ST) + \ell + \frac{(n-\ell)\lambda}{n} - ST - \ell T \right) - \ell T, \\
&= (c-1) \left(\alpha\lambda + \ell + \frac{(n-\ell)\lambda}{n} - T((\alpha+1)S + \ell) \right) - \ell T, \\
&= (c-1) \left((2n-1)\lambda + \ell + \frac{(n-\ell)\lambda}{n} - T(2nS + \frac{c}{c-1}\ell) \right).
\end{aligned}$$

Note that if we examine the backoff counters of nonempty users at the start of any step (i.e. before we add any arrivals), there can be at most one user with a backoff counter equal to 0 (i.e. a user that will definitely send). This is due to the fact that there can be only one successful user at the end of the previous step. Showing that the drift is nonnegative will be divided into two cases according to the existence of a nonempty user with a counter equal to 0 at the start of the current step.

Case 1. $q_1 \geq 1, \dots, q_{\ell} \geq 1$ **and** $q_{\ell+1} = 0, \dots, q_n = 0$; $b_1 \geq 1, \dots, b_{\ell} \geq 1$; $1 \leq \ell \leq n$

In this case there are ℓ active users and none has a backoff counter equal to 0 at the start of the step. Define $\epsilon_i = \frac{p_i}{1-p_i}$, and note that $1 + \epsilon_i = \frac{1}{1-p_i}$ and that

$0 < \epsilon_i < 1$ because in this case $0 < p_i < 1$. Next, define

$$R = \prod_{i=1}^n (1 + \epsilon_i).$$

Note that $R = 1/T$, and $S = \sum_{i=1}^n \epsilon_i$. For any $\lambda \geq \frac{1}{2} + \frac{1}{4n-2}$, we will show that in this case

$$E[f(X_{t+1}) - f(X_t)] \geq (c-1) \left(n + \ell + \frac{(n-\ell)\lambda}{n} - T(2nS + \frac{c}{c-1}\ell) \right) \geq 0.$$

Because $c \geq 2$, the last inequality is satisfied when

$$n + \ell \geq T(2nS + 2\ell),$$

which can be rewritten as

$$R(n + \ell) \geq 2nS + 2\ell. \tag{4.2}$$

To verify that Inequality (4.2) is true, we will use the following lemma which is due to Håstad, Leighton and Rogoff (Lemma 5.2 in [27]).

Lemma 4.3 *If $0 \leq \epsilon_i \leq 1$ for $1 \leq i \leq m$, then*

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

and

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

Note that the first inequality in the lemma shows that

$$R \geq 1 + S,$$

and the second show that

$$R \geq 2S.$$

Hence,

$$\begin{aligned} R(n + \ell) &\geq R(n - \ell) + 2R\ell \\ &\geq 2nS - 2\ell S + 2(1 + S)\ell \\ &= 2nS + 2\ell. \end{aligned}$$

This shows that Inequality (4.2) is satisfied.

Case 2. $q_1 \geq 1, \dots, q_{\ell+1} \geq 1$ **and** $q_{\ell+2} = 0, \dots, q_n = 0$; $b_1 \geq 1, \dots, b_\ell \geq 1$; $b_{\ell+1} = 0$; $0 \leq \ell < n$

In this case (and at the beginning of the step before new arrivals are added to the queues) there is a single user (i.e. user $\ell + 1$) with a nonempty queue and a backoff counter equal to 0.² In this case, if any user other than user $\ell + 1$ sends, there will be a collision. Therefore, only this user can send successfully when all of the others stay silent. The probability that this happens (and user $\ell + 1$ succeeds) is

$$W = \prod_{i=1}^{\ell} (1 - c^{-b_i}) \prod_{i=\ell+2}^n (1 - \lambda_i).$$

We calculate the drift in this case as

$$\begin{aligned} & E[f(X_{t+1}) - f(X_t)] \\ &= (c-1)\alpha \sum_{i=1}^n (E[I_{a,i,t}] - E[I_{s,i,t}]) + \sum_{i=1}^n (E[c^{b_i(t+1)}] - c^{b_i(t)}) \\ &= (c-1)\alpha(\lambda - W) + \sum_{i=1}^n ((c^{b_i(t)+1} - c^{b_i(t)})\sigma_i - (c^{b_i(t)} - 1)\rho_i) \\ &= (c-1)\alpha(\lambda - W) + \sum_{i=1}^{\ell} (c-1) + (c-1)(1 - W) + (c-1) \sum_{i=\ell+2}^n \lambda_i \\ &= (c-1) \left(\alpha(\lambda - W) + \ell + (1 - W) + \sum_{i=\ell+2}^n \lambda_i \right). \end{aligned}$$

Thus, the expected change in the potential is

$$E[f(X_{t+1}) - f(X_t)] = (c-1) \left(\alpha(\lambda - W) + \ell + (1 - W) + \sum_{i=\ell+2}^n \lambda_i \right).$$

Since $c > 1$, this is nonnegative when

$$\alpha\lambda + \ell + 1 + \sum_{i=\ell+2}^n \lambda_i \geq (\alpha + 1)W.$$

Since the arrival rates at each user are symmetric, $\lambda_i \geq \lambda/n$. Substituting constants into the inequality, it suffices to show that the function

$$g(\ell) = (2n-1)\lambda + \ell + 1 + (n-\ell-1)\frac{\lambda}{n} - 2n \left(1 - \frac{\lambda}{n} \right)^{n-\ell-1}$$

²From the case description, we know that users $1 \dots \ell$ have nonempty queues and $b_i > 0$, user $\ell + 1$ has a nonempty queue and $b_{\ell+1} = 0$ and the users $\ell + 2 \dots n$ have empty queues at the start of the step.

is nonnegative for all possible values of ℓ .

To do this we will show that

1. $g(0)$ is nonnegative,
2. $g(n-1)$ is nonnegative, and
3. $\frac{\partial^2}{\partial \ell^2} g(\ell) < 0$ for $0 \leq \ell \leq n-1$.

To prove that $g(n-1) \geq 0$, note that

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

is nonnegative when

$$\lambda \geq \frac{n}{2n-1} = \frac{1}{2} + \frac{1}{4n-2}.$$

Next, we show that $\frac{\partial^2}{\partial \ell^2} g(\ell) < 0$ for $0 \leq \ell \leq n-1$. This is clearly true since the second derivative is

$$\frac{\partial^2}{\partial \ell^2} g(\ell) = -2n \left(1 - \frac{\lambda}{n}\right)^{n-\ell-1} \left(\ln \left(1 - \frac{\lambda}{n}\right)\right)^2 < 0.$$

Finally, we must show that

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

This is a bit tricky, and we will prove it using the following calculation due to Håstad et. al. [27] which we copy here with its proof.

Proposition 4.4 (Håstad, Leighton, and Rogoff [27]) *Suppose that λ_0 is the solution of the equation $\lambda_0 = e^{-\lambda_0}$. Then for any $\lambda \geq \lambda_0 + \frac{1}{4n-2}$,*

$$\left(1 - \frac{\lambda}{n}\right)^{n-1} \leq \lambda.$$

Proof: First note that

$$1 - \frac{\lambda}{n} = e^{-\frac{\lambda}{n} - \frac{\lambda^2}{2n^2} - \frac{\lambda^3}{3n^3} - \dots}$$

and thus

$$\begin{aligned}
\left(1 - \frac{\lambda}{n}\right)^{n-1} &= e^{-\lambda + \frac{\lambda}{n} - \frac{\lambda^2}{2n} + \frac{\lambda^2}{2n^2} - \frac{\lambda^3}{3n^2} + \frac{\lambda^3}{3n^3} - \dots} \\
&= e^{-\lambda + \frac{\lambda(2-\lambda)}{2n} + \frac{\lambda^2(3-2\lambda)}{6n^2} + \frac{\lambda^3(4-3\lambda)}{12n^3} + \dots} \\
&< e^{-\lambda + \frac{1}{2n} + \frac{1}{6n^2} + \frac{1}{12n^3} + \dots} && \text{since } \lambda < 1, \\
&< e^{-\lambda + \frac{1}{2n} + \frac{1}{6n^2} \left(\frac{1}{1-1/n}\right)} \\
&\leq e^{-\lambda + \frac{1}{2n-1}} \\
&= e^{-(\lambda - \frac{1}{4n-2}) + \frac{1}{4n-2}} \\
&\leq \left(\lambda - \frac{1}{4n-2}\right) e^{\frac{1}{4n-2}} && \text{since } \lambda - \frac{1}{4n-2} \geq \lambda_0 \\
&&& \text{and } e^{-\lambda'} \leq \lambda' \text{ for all } \lambda' \geq \lambda_0 \\
&\leq \left(\lambda - \frac{1}{4n-2}\right) \frac{1}{1 - \frac{1}{4n-2}} && \text{since } e^x \leq \frac{1}{1-x} \text{ for } 0 \leq x < 1, \\
&\leq \lambda.
\end{aligned}$$

□

Using Proposition 4.4 we conclude that

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

Thus we have shown that $g(\ell) \geq 0$ for all $0 \leq \ell \leq n-1$, which proves that $E[f(X_{t+1}) - f(X_t)] \geq 0$ in this case.

These two cases prove that Condition 2 is true, and thus using Theorem 4.1 the protocol is unstable for any $\lambda > \lambda_0 + \frac{1}{4n-2}$. □

4.2 Instability of the Two-User Binary Exponential Backoff

Consider the binary exponential backoff protocol running on a system of two users sharing a multiple-access channel. By working with only two users, our aim is to seek some exact answers for the stability region of this protocol. However, we will see in this section that reducing the number of users does not necessarily make the analysis trivial. It seems frustrating that even though limiting the number of users to two might simplify the problem, there are still

many open questions about the stability region. Yet, at the same time, this simplified problem shows that finding answers about the stability of backoff protocols can be very a non-trivial and interesting task.

Results for this case were first obtained by Goodman, Greenberg, Madras, and March [25] who showed that binary exponential backoff is positive recurrent when $\lambda < 0.3$ given that the arrival rates are symmetric. For non-symmetric arrivals, they showed that the system is unstable when either λ_1 or λ_2 is at least $1/2$. The same authors obtained some conditions on the two arrival rates under which the protocol is stable when these rates are non-symmetric. The effort this paper had to go through to obtain this result is another illustration of the complexity of analysing backoff protocols even in a simple situation such as the two-user case.

Håstad, Leighton and Rogoff [27] later showed (as a special case of their instability result) that when the two arrival rates are symmetric, the system is unstable when $\lambda > 0.733667$. In this section we will show that for symmetric arrivals, binary exponential backoff is unstable for $\lambda > 2/3$. We will also give some conditions under which the protocol is unstable for non-symmetric arrivals. These results (including ours) are illustrated in Figure 4.1 which summarises what we know about the two-user case of binary exponential backoff.

The main result of this section is presented by the following theorem.

Theorem 4.5 *Suppose that two users are sharing a multiple-access channel using the binary exponential backoff protocol with mean arrival rates λ_1 and λ_2 . Then the system is unstable when any of the following inequalities are satisfied*

1. $\lambda_1 \geq 1/5$ and $\lambda_2 \geq 1/5$ and $\lambda_1 + \lambda_2 \geq 2/3$; or
2. $\lambda_1 < 1/5$ and $8\lambda_1 + 3\lambda_2 \geq 3$; or
3. $\lambda_2 < 1/5$ and $3\lambda_1 + 8\lambda_2 \geq 3$.

The instability region shown by this result is illustrated in Figure 4.1 which also compares it with previous results by other authors. When the ar-




Figure 4.1: The stability and instability regions for the two-users binary exponential backoff.

rival rates into the two stations are symmetric, we get the following simplified corollary.

Corollary 4.6 *Any system with two users running the binary exponential back-off protocol with mean arrival rates $\lambda_1 = \lambda_2 = \lambda/2$ is unstable for any $\lambda > 2/3$.*

The system will be analysed using a Markov chain X_t and an associated potential function. Each state in this chain will consist of a quadruple (q_1, q_2, b_1, b_2) specifying the current queue sizes and the values of the backoff counters. The proof will use Theorem 4.1 to show that this Markov chain is non-positive recurrent. Note that this chain is time-homogeneous, aperiodic, irreducible and has a countable state space $\mathbf{X}$.

We begin the analysis of the chain by choosing a suitable potential function. For each state X_t define the potential as

$$f(X_t) = 3(q_1(t) + q_2(t)) + 2^{b_1(t)} + 2^{b_2(t)}.$$

4.2.1 Bounded Expected Jumps

The following lemma shows that Condition 3 in Theorem 4.1 is satisfied.

Lemma 4.7 *Given any state x , there exists a finite constant d such that*

$$E[|f(X_{t+1}) - f(X_t)| \mid X_t = x] \leq d.$$

Proof: Because of the absolute value on the change in potential, we will divide the proof into two cases depending on whether the change is negative or not. Let E denote the event that $f(X_{t+1}) - f(X_t) \geq 0$. Thus, E is true when the change is nonnegative.

Case E :

In this case, the change is nonnegative. From the analysis of a single back-off counter in a blocked channel (i.e. where no user succeeds), we know that $E[2^{b_i(t+1)}] = 2^{b_i(t)} + 1$. Thus,

$$E[|f(X_{t+1}) - f(X_t)| \mid X_t = x] \leq 3(\lambda_1 + \lambda_2) + 1 + 1 \leq 5. \quad (4.3)$$

Case $\overline{E}$:

The change in potential in this case is negative. This can be caused by one event only, namely that a user succeeds in sending. This will result in a drop by 3 (from a message being removed from the queue) and $2^{b_i} - 1$ (from backoff counter i being reset to 0). The drop from the backoff counter can be very large. However, this happens with probability

$$\Pr(\text{User } i \text{ succeeds}) \leq 2^{-b_i}.$$

Thus, the expected drop in the potential is at most

$$2^{-b_i}(3 + 2^{b_i}) \leq 3 + 1. \quad (4.4)$$

From Equation (4.3) and Equation (4.4), we conclude that

$$\begin{aligned} E[|f(X_{t+1}) - f(X_t)| \mid X_t = x] &\leq E[f(X_{t+1}) - f(X_t) \mid X_t = x \wedge E] \\ &\quad + E[f(X_{t+1}) - f(X_t) \mid X_t = x \wedge \overline{E}] \\ &\leq 5 + 4 \\ &\leq d. \end{aligned}$$

□

4.2.2 Non-negative drift

Let the constant V_0 in the statement of Theorem 4.1 be equal to 2. This establishes Condition 1 of Theorem 4.1 since the set of states x with $f(x) \leq V_0$ includes the empty state $(0, 0, 0, 0)$ only. The following lemma establishes Condition 2 of Theorem 4.1.

Lemma 4.8 *Given any state x with $f(x) > 2$, and assuming that the following conditions on the mean arrival rates are true:*

- (i) $\lambda_1 + \lambda_2 \geq 2/3$.
- (ii) $3\lambda_1 + 8\lambda_2 \geq 3$.

$$(iii) \quad 8\lambda_1 + 3\lambda_2 \geq 3.$$

Then

$$E[f(X_{t+1}) - f(X_t) | X_t = x] \geq 0.$$

Proof:

The analysis is divided into five cases depending on whether a queue or a backoff counter is zero or not, and the expected change in potential is calculated over a single step in every case. We will exclude the empty state $(0, 0, 0, 0)$ by requiring that the nonnegative expected change is observed in states x with potential $f(x) > 2$. This was done by setting $V_0 = 2$. Inequality (i) on the arrival rates which requires that

$$\lambda_1 + \lambda_2 \geq \frac{2}{3} \tag{4.5}$$

will be sufficient to show that the drift in most cases is nonnegative. However, in the first two cases another weaker condition will be derived. This will give us a slightly higher rate for instability when $\lambda_1 < 1/5$ or $\lambda_2 < 1/5$.

Case 1a : States $x = (q_1, 0, 0, 0)$ where $q_1 \geq 1$

In this case there is only one nonempty queue with backoff counter equal to zero. Thus, queue 1 will send with probability one in this step. Collisions can only occur when a new message arrives at user 2 during this step. The analysis of this case is simplified by the fact that both stations send with a known probability. Thus, we can calculate the exact value of the expected change in the potential.

$$E[f(X_{t+1}) - f(X_t) | X_t = x] = 3(\lambda_1 + \lambda_2) - 3(1 - \lambda_2) + \lambda_2(1 + 1) - 0.$$

By simplifying the terms, this is nonnegative when the following inequality is true

$$3\lambda_1 + 8\lambda_2 \geq 3. \tag{4.6}$$

This is Condition (ii) in the statement of the lemma.

Case 1b : States $x = (0, q_2, 0, 0)$ **where** $q_2 \geq 1$

This case is very similar to the previous case, but now user 2 has the nonempty queue. The expected change in a single step is

$$E[f(X_{t+1}) - f(X_t) | X_t = x] = 3(\lambda_1 + \lambda_2) - 3(1 - \lambda_1) + \lambda_1(1 + 1) - 0.$$

Again, by simplifying we get the following condition for this to be nonnegative

$$8\lambda_1 + 3\lambda_2 \geq 3. \quad (4.7)$$

This shows that Condition (iii) in the statement of the lemma is true.

Case 2a : States $x = (q_1, 0, b_1, 0)$ **where** $q_1 \geq 1, b_1 \geq 1$

In these states, user 1 has a nonempty queue. However, its backoff counter is larger than 0. Therefore, it will not definitely send in the next step. The probability that user 1 sends is 2^{-b_1} while the probability that an arrival comes to user 2 and is sent is λ_2 .

Suppose that Inequality (4.5) (which is equivalent to Condition (i)) is true. We will show that the expected change in the potential is nonnegative.

$$\begin{aligned} E[f(X_{t+1}) - f(X_t) | X_t = x] &= 3(\lambda_1 + \lambda_2) - 3 \left(2^{-b_1}(1 - \lambda_2) + \lambda_2(1 - 2^{-b_1}) \right) \\ &\quad + \lambda_2(2^{-b_1} + 1) - (1 - \lambda_2)(1 - 2^{-b_1}) \\ &= -1 - 2^{1-b_1} + 3\lambda_1 + 2\lambda_2 + 3 \cdot 2^{1-b_1} \lambda_2 \\ &= -1 - 2^{1-b_1} + 3(\lambda_1 + \lambda_2) - \lambda_2 + 3 \cdot 2^{1-b_1} \lambda_2 \\ &\geq -1 - 2^{1-b_1} + 2\left(\frac{2}{3}\right) + \lambda_1 + 3 \cdot 2^{1-b_1} \lambda_2 \\ &= \frac{1}{3} - 2^{1-b_1} + \lambda_1 + 3 \cdot 2^{1-b_1} \lambda_2 \\ &= 2^{-b_1} \times \left(\frac{2^{b_1}}{3} - 2 + 2^{b_1} \lambda_1 + 6\lambda_2 \right) \\ &\geq 2^{-b_1} \times \left(\frac{2}{3} - 2 + 2\lambda_1 + 6\lambda_2 \right) \\ &\geq 2^{-b_1} \times \left(\frac{2}{3} - 2 + 2\left(\frac{2}{3}\right) + 4\lambda_2 \right) \\ &= 2^{-b_1} 4\lambda_2 \\ &\geq 0. \end{aligned}$$

Case 2b : States $x = (0, q_2, 0, b_2)$ where $q_2 \geq 1, b_2 \geq 1$

These states are very similar to the states in Case 2a, and by symmetry we can show that the expected change in the potential is nonnegative.

Case 3: $x = (q_1, q_2, b_1, b_2)$ where $q_1 \geq 1, q_2 \geq 1, b_1 \geq 0$, and $b_2 \geq 0$

Both users have nonempty queues. Therefore, the arrivals have no immediate impact on the collisions. The only effect of arrivals is to increase the queue sizes. Assuming that Inequality (4.5) (i.e. Condition (i)) is true, the expected change in a single step is

$$\begin{aligned}
E[f(X_{t+1}) - f(X_t) | X_t = x] &= 3(\lambda_1 + \lambda_2) - 3(2^{-b_1}(1 - 2^{-b_2}) + 2^{-b_2}(1 - 2^{-b_1})) \\
&\quad + 2^{-b_1} + 2^{-b_2} - 2(1 - 2^{-b_1})(1 - 2^{-b_2}) \\
&\geq 3\left(\frac{2}{3}\right) - 3(2^{-b_1}(1 - 2^{-b_2}) + 2^{-b_2}(1 - 2^{-b_1})) \\
&\quad + 2^{-b_1} + 2^{-b_2} - 2(1 - 2^{-b_1})(1 - 2^{-b_2}) \\
&= 2^{2-b_1-b_2} \\
&\geq 0.
\end{aligned}$$

□

Now we are ready to prove the main theorem in this section.

Proof of Theorem 4.5: In order to show that the drift is nonnegative, we must ensure that inequalities (i), (ii), and (iii) are all satisfied. Although when the two arrival rates are close to each other the system is unstable for $\lambda_1 + \lambda_2 \geq 2/3$, we must use other conditions when one arrival rate is much smaller than the other. The proof is divided into the three cases taken from the statement of Theorem 4.5.

Case 1. $\lambda_1 \geq 1/5, \lambda_2 \geq 1/5$, and $\lambda_1 + \lambda_2 \geq 2/3$

Consider first the case where both arrival rates are at least $1/5$, and their sum is at least $2/3$. Thus, we are given Condition (i) and we must show that Conditions (ii) and (iii) are satisfied. These can be easily shown by the following. For

Condition (ii), note that using the fact that $\lambda_1 + \lambda_2 \geq 2/3$ and $\lambda_2 \geq 1/5$, we can deduce

$$\begin{aligned} 3\lambda_1 + 8\lambda_2 &= 3(\lambda_1 + \lambda_2) + 5\lambda_2 \\ &\geq 3(2/3) + 5(1/5) \\ &= 3. \end{aligned}$$

Similarly, using $\lambda_1 + \lambda_2 \geq 3$ and $\lambda_1 \geq 1/5$, the following shows that Condition (iii) is satisfied.

$$\begin{aligned} 8\lambda_1 + 3\lambda_2 &= 5\lambda_1 + 3(\lambda_1 + \lambda_2) \\ &\geq 5(1/5) + 3(2/3) \\ &= 3 \end{aligned}$$

Thus in this case all of the conditions of Lemma 4.8 are met and the drift is nonnegative.

Case 2. $\lambda_1 < 1/5$ and $8\lambda_1 + 3\lambda_2 \geq 3$

Consider case number 2 in the statement of Theorem 4.5. In this case the first user has a low arrival rate, namely $\lambda_1 < 1/5$. We must show that in this case, if $8\lambda_1 + 3\lambda_2 \geq 3$ (i.e. Condition (iii)) is satisfied then both (i) and (ii) are true. Assume that $\lambda_1 < 1/5$ and note that $8\lambda_1 + 3\lambda_2 \geq 3$ can be rewritten as

$$\begin{aligned} \lambda_1 + \lambda_2 &\geq 1 - \frac{5}{3}\lambda_1 \\ &\geq 1 - \frac{1}{3} \\ &= 2/3 \end{aligned}$$

which is equal to Condition (i).

To show that Condition (ii) is satisfied, note that since we have shown that

$$\lambda_1 + \lambda_2 \geq 2/3,$$

$$\begin{aligned} 3\lambda_1 + 8\lambda_2 &\geq 3\lambda_1 + 8(2/3 - \lambda_1) \\ &\geq 8 \cdot \frac{2}{3} - 5 \cdot \lambda_1 \\ &\geq \frac{16}{3} - 5 \cdot \frac{1}{5} \\ &= \frac{13}{3} \\ &\geq 3. \end{aligned}$$

Case 3. $\lambda_2 < 1/5$ **and** $3\lambda_1 + 8\lambda_2 \geq 3$

Again, by symmetry, this case is identical to the previous case.

In any of the three cases we have shown that Conditions (i), (ii), and (iii) are all satisfied. Using Lemma 4.8 this shows that statement 2 of Theorem 4.1 is satisfied (i.e. the chain has nonnegative drift under the potential function f). Lemma 4.7 shows that statement 3 of Theorem 4.1 is also satisfied.

Furthermore, since $V_0 = 2$ then statement 1 is satisfied. Thus, using Theorem 4.1 we conclude that the Markov chain X_t is not positive recurrent and has no stationary distribution. $\square$

Chapter 5

Age-based Protocols in the Finite Model

While backoff protocols have gained popularity after their initial use in the Ethernet, only a small number of other acknowledgement-based protocols have been studied in the literature. These include age-based protocols, which were first studied for the infinite users model by MacPhee [38] in his PhD thesis and in the paper by Kelly and MacPhee [35] (see also the survey by Kelly [34]). For the finite users model, no stability study has been conducted for age-based protocols until now.

Age-based and backoff protocols are both acknowledgement-based, and thus share the nice properties outlined in Chapter 1. The main difference between the two classes of protocols is in the way the probability of retransmission is calculated. In backoff protocols, the transmission probability is a function of the number of collisions encountered by a message. While in age-based protocols it is a function of the age of the message (i.e. the time it has been waiting to be sent). For the finite model, the age is defined as the time a message has been waiting at the head of the queue. We only consider the age of a message after it reaches the head of a queue because in the finite model only these messages participate in the contention process.

In the next section we will present a protocol which is, in a sense, similar to binary exponential backoff. However, while binary exponential backoff sends with an expected probability $E[p_i] = \Theta(1/a_i)$ for a message that has been waiting for a_i steps; the age-based protocol sends with probability $p_i = \Theta(1/a_i)$. We will show that this protocol is essentially stable for a slightly higher λ than the rate known to be sufficient for exponential backoff to be positive recurrent. Furthermore, the form of stability used for the age-based protocol presented below will be stronger than positive recurrence, in the sense that we require that $E[T_{ret}] < \infty$ and that the expected load of the system in the stationary distribution (i.e. $E_\pi[L(X_t)]$) is also bounded. To prove this strong form of stability we will use the results of Chapter 2. Note that the results of this chapter are asymptotic in the number of users n .

5.1 Strong Stability of an Age-based Protocol Similar to Binary Exponential Backoff

The model assumptions for this protocol are basically the same as those presented in Section 3.1 for c -ary exponential backoff. We will again consider a system of n users sharing a multiple-access channel. Every user has a queue where it stores arriving messages. In each step, a single arrival into user i is added to its queue with probability λ_i , and none with probability $1 - \lambda_i$. Suppose that arrivals are symmetric, so that the rate for every user i is $\lambda_i = \lambda/n$. Time is divided into discrete steps (i.e. we are working in the so called slotted model) and a single message is small enough so that it can be sent in a single time step.

The queue size of the i th user at the beginning of the t 'th time step is denoted by $q_i(t)$. Instead of backoff counters, users maintain age counters. Let $a_i(t)$ denote the age of the message stored at the head of the queue of user i . (The age of a message is equal to the number of time steps since it arrived at the head of the queue, and empty users have counters equal to 0.) Thus a

message born at step t (i.e. arriving at the head of queue i at step t) has age $a_i(t) = 0$. If the same message has not been sent at step $t + T$, then its age will be $a_i(t + T) = T$.

The age-based protocol works as follows. At each step t , user i will send with probability $\frac{4}{a_i(t)+4}$ if the head of its queue is not empty. Since this is a multiple-access channel, if any one else sends at the same time step, a collision is detected and none of the transmissions are successful. However, if user i is the only one sending, the transmission is considered successful, and the message is removed from the queue of user i .

To analyse this protocol, we will again model it by a Markov chain X_t . The chain has a countable state space $\mathbf{X} = (\mathbb{Z}^+)^{2n}$. A state of the system at time t will be represented by the following $2n$ -tuple,

$$X_t = (q_1(t), \dots, q_n(t), a_1(t), \dots, a_n(t)).$$

From each state, the probability of moving to another state is determined by the arrival process and the send probabilities defined above for the protocol. Similar to the chains considered before, this chain is time-homogeneous, irreducible, and aperiodic.

Let us again define the system load at time t as $L(X_t) = \sum_{i=1}^n q_i(t)$. Each state X_t in the Markov chain is associated with a positive potential using the following function.

$$f(X_t) = \alpha n^{3/2} \sum_{i=1}^n q_i(t) + \sum_{i=1}^n a_i(t) + 1.$$

Note that this function is an upper bound on the system load. Using this potential function we will show the following result.

Theorem 5.1 *Suppose that an age-based protocol is running on a system of n users, and that each user with a nonempty queue sends with probability $p_i = 4/(a_i + 4)$, where a_i is the age of the message at the head of queue i . Furthermore, suppose that arrivals at user i are Bernoulli distributed with probability $\lambda_i = \lambda/n$ of receiving an arrival at each step. Then there exists a constant α such that*

the system is strongly stable (in the sense that $E[T_{ret}] < \infty$ and there exists a unique stationary distribution π such that $E_\pi[L(X_t)] < \infty$) for sufficiently large n and any $\lambda \leq 1/(\alpha n^{3/4})$.

We will show that Conditions (C1), (C2), and (V2) in Lemma 2.8 are satisfied by the Markov chain X_t . The proof of Condition (C2) follows the same argument used to show that c -ary exponential backoff is positive recurrent. However, the proof for the age-based protocol (i.e. for Condition (C2)) is simpler and shorter than the backoff proof because we do not need to use probabilistic bounds to estimate the probability of an attempt (like we did in Lemma 3.2). This was necessary with backoff because of the stochastic nature of the backoff counters and the fact that after T steps from the start state we do not know how many collisions have been encountered by user i . This is not the case in an age-based protocol since after T steps, if the message has not been sent successfully then we know that the age counter has increased by exactly T . The same reason is responsible for making it possible to show condition (C1) is true ¹, and for the stronger and slightly better stability rate bound for the age-based protocol than c -ary exponential backoff (i.e. we can show that the expected load in the stationary distribution is bounded for the age-based protocol, and $\lambda \leq 1/(\alpha n^{3/4})$ for age-based instead of $1/(\alpha n^{3/4+\epsilon})$ for exponential backoff).

Let us define the function $k(x) \geq 1$ to be an integer function which takes its values from the set $\{1, \dots, K\}$, where $1 \leq K < \infty$. This function will determine the number of steps needed in each case to show negative drift. Let us also set here the values of the constants from Lemma 2.8 as $\beta = 17$, $\sigma = 2.75$, $\alpha = e^{20}\beta$, $K = 3n$, and $V_0 = 64\alpha^2 n^5 K^2$. Let C be the set of all states s with potential $f(s) \leq V_0$. Note that C is a finite set. The function $V(s)$ will be set to denote $f^2(s)$ for all states $s \in \mathbf{X}$.

We start the proof by showing that the two simpler Conditions (C1) and (V2) are satisfied. Then after showing that Condition (C2) is also satisfied, we

¹To be more specific, the main reason that enabled us to show condition (C1) for the potential function used for the age-based and not the backoff protocol is the fact that the upward jumps for the former are bounded.

will present the proof of Theorem 5.1.

5.1.1 Conditions (C1) and (V2)

The following lemma establishes Condition (C1) for the age-based protocol defined above.

Lemma 5.2 *Given the chain X_t representing the evolution of the age-based protocol, and given the potential function f defined above, then Condition (C1) in Lemma 2.8 is satisfied.*

Proof: For convenience, let us restate Condition (C1). Remember that the function $V(s)$ was set to $f^2(s)$, and given any state $s \in \mathbf{X} - C$ such that $s' = X_{t+k(s)}$, we must show

$$(C1) \quad E[(f(s') - f(s))^2] \leq (2\sigma - 1) k(s) \cdot f(s).$$

To prove this condition we start with some notation. Let

$$Q(X_t) = \alpha n^{3/2} \sum_{i=1}^n q_i(t)$$

and

$$A(X_t) = \sum_{i=1}^n a_i(t).$$

Note that the potential function can be written as $f(X_t) = Q(X_t) + A(X_t) + 1$. The following non-negative random variables will also simplify the analysis. Let

$$\begin{aligned} Q^+ &= \max\{0, Q(s') - Q(s)\}, \\ Q^- &= \max\{0, -(Q(s') - Q(s))\}, \\ A^+ &= \max\{0, A(s') - A(s)\}, \\ A^- &= \max\{0, -(A(s') - A(s))\}. \end{aligned}$$

There is a clear relationship between these nonnegative random variables and the drift, that is obscured by the max function. To establish this relation we will need the following simple proposition.

Proposition 5.3 *For any two real numbers x and y ,*

$$x - y = \max\{0, x - y\} - \max\{0, y - x\}.$$

Proof: There are three cases. If $x = y$ then both sides are 0 and the equality holds. If $x > y$ then the right hand side is $\max\{0, x - y\} - 0$, which is equal to the left hand side $x - y$. On the otherhand, if $x < y$ then the right hand side is equal to $0 - \max\{y - x, 0\}$, which is also equal to the left hand side. $\square$

Using this proposition, we can show that $Q(s') - Q(s) = Q^+ - Q^-$ and $A(s') - A(s) = A^+ - A^-$, and thus

$$f(s') - f(s) = Q^+ - Q^- + A^+ - A^-.$$

The expected square of the change in the potential can now be written as,

$$\begin{aligned} E[(f(s') - f(s))^2] &= E[(Q^+ - Q^- + A^+ - A^-)^2] \\ &\leq E[(Q^+)^2 + (A^+)^2 + (Q^-)^2 + (A^-)^2 + 2Q^-A^- + 2Q^+A^+]. \end{aligned}$$

Most of these random variables are trivially bounded from above. Notice that $Q^+ \leq \alpha n^{3/2} \cdot nK$ because we can have at most n new arrivals into the queues at each step, and there are $k(s) \leq K$ steps. Also, $Q^- \leq \alpha n^{3/2} K$ since at most one message is sent successfully at each step, and there are at most K steps. For the age counter part of the potential, $A^+ \leq nK$, since the counter of each user can increase by at most 1, and there are at most n users. Using these bounds,

$$\begin{aligned} E[(f(s') - f(s))^2] &\leq (\alpha n^{3/2} \cdot nK)^2 + (nK)^2 + (\alpha n^{3/2} K)^2 + E[(A^-)^2] \\ &\quad + 2\alpha n^{3/2} K E[A^-] + 2\alpha n^{3/2} \cdot (nK)^2. \end{aligned} \tag{5.1}$$

Next, with slightly more effort we will bound $E[A^-]$ and $E[(A^-)^2]$. For each user i , let

$$A_i^- = \max\{0, -(a_i(t + k(s)) - a_i(t))\}$$

be the drop in the value of the age counter of user i , and note that

$$A^- \leq \sum_{i=1}^n A_i^-.$$

Let $t' = t + k(X_t)$. The bound on $E[A^-]$ is obtained as follows,

$$\begin{aligned}
E[A^-] &\leq E\left[\sum_{i=1}^n A_i^-\right] \\
&= \sum_{i=1}^n E[A_i^-] \\
&= \sum_{i=1}^n E[\max\{0, -(a_i(t') - a_i(t))\}] \\
&= \sum_{i=1}^n \sum_{x: x \geq 0} x \cdot \Pr(\max\{0, -(a_i(t') - a_i(t))\} = x) \\
&= \sum_{i=1}^n \left(\sum_{x: x > a_i(t)} 0 \right. \\
&\quad \left. + \sum_{x: 0 < x \leq a_i(t)} x \cdot \Pr(a_i(t) - a_i(t') = x) \right) \tag{5.2}
\end{aligned}$$

$$\begin{aligned}
&\leq \sum_{i=1}^n \sum_{x: 0 < x \leq a_i(t)} a_i(t) \cdot \Pr(a_i(t) - a_i(t') = x) \tag{5.3} \\
&= \sum_{i=1}^n a_i(t) \Pr(a_i(t') < a_i(t))
\end{aligned}$$

The zero sum in Equation (5.2) is caused by the fact that the age counter cannot drop by more than what it started with (i.e. $a_i(t)$). The inequality in Equation (5.3) is due to the same reason.

Furthermore, notice that

$$\Pr(a_i(t) < a_i(t')) \leq \frac{4}{a_i(t) + 4} \cdot K$$

since in order for the counter to decrease, the user must attempt to send and succeed, and there are at most K chances of doing so. This leads to the following bound on the expected drop in the age counters

$$E[A^-] \leq \sum_{i=1}^n a_i(t) \frac{4}{a_i(t) + 4} \cdot K \leq 4nK.$$

To calculate the remaining bound on $E[(A^-)^2]$, we will need to establish the following proposition.

Proposition 5.4 *Given any sequence $x_1, \dots, x_n$ of real numbers, the following inequality is true*

$$\left(\max\{0, \sum_{i=1}^n x_i\} \right)^2 \leq \sum_{i=1}^n (\max\{0, x_i\})^2 + 2 \sum_{i,j \wedge i < j} \max\{0, x_i\} \cdot \max\{0, x_j\}.$$

Proof: Assume first that $\sum_{i=1}^n x_i \leq 0$, then the left hand side is equal to 0, and the right hand side is clearly nonnegative. Hence, the inequality is true in this case.

On the other hand, assume that $\sum_{i=1}^n x_i > 0$. In this case,

$$\left(\max\{0, \sum_{i=1}^n x_i\} \right)^2 = \left(\sum_{i=1}^n x_i \right)^2.$$

Let m be the number of nonnegative numbers among the sequence $x_1, \dots, x_n$. Assume without loss of generality that the first m numbers in this sequence are nonnegative. Let $X_+ = \sum_{i=1}^m \max\{0, x_i\}$, and $X_- = \sum_{i=1}^n \max\{0, -x_i\}$. Notice that both are nonnegative quantities and that $\sum_{i=1}^n x_i = X_+ - X_-$. Then,

$$\begin{aligned} \left(\sum_{i=1}^n x_i \right)^2 &= (X_+ - X_-)^2 \\ &= X_+^2 + X_-^2 - 2X_+ \cdot X_- \\ &\leq X_+^2 \end{aligned}$$

since in this case $X_+ > X_-$.

Thus we conclude that in this case (i.e. when $X_+ > X_-$),

$$\begin{aligned} \left(\max\{0, \sum_{i=1}^n x_i\} \right)^2 &\leq \left(\sum_{i=1}^m x_i \right)^2 \\ &= \sum_{i=1}^m x_i^2 + 2 \sum_{i,j \in \{1, \dots, m\} \wedge i < j} x_i \cdot x_j \\ &= \sum_{i=1}^m (\max\{0, x_i\})^2 + 2 \sum_{i,j \wedge i < j} \max\{0, x_i\} \cdot \max\{0, x_j\}. \end{aligned}$$

This concludes the proof of the proposition. $\square$

By Proposition 5.4,

$$\begin{aligned}
E[(A^-)^2] &\leq E\left[\sum_{i=1}^n (A_i^-)^2 + 2 \sum_{i,j \wedge i < j} A_i^- \cdot A_j^-\right] \\
&= \sum_{i=1}^n E[(A_i^-)^2] + 2 \sum_{i,j \wedge i < j} E[A_i^- \cdot A_j^-] \tag{5.4}
\end{aligned}$$

We bound the term $E[A_i^- \cdot A_j^-]$ as follows.

$$\begin{aligned}
E[A_i^- \cdot A_j^-] &= \sum_{x: 1 \leq x \leq a_i(t)a_j(t)} x \cdot \Pr(A_i^- \cdot A_j^- = x) \\
&\leq \sum_{x: 1 \leq x \leq a_i(t)a_j(t)} a_i(t)a_j(t) \Pr(A_i^- \cdot A_j^- = x) \\
&= a_i(t)a_j(t) \sum_{x: 1 \leq x \leq a_i(t)a_j(t)} \Pr(A_i^- \cdot A_j^- = x) \\
&= a_i(t)a_j(t) \Pr(1 \leq A_i^- \cdot A_j^- \leq a_i(t)a_j(t))
\end{aligned}$$

We only consider values of x where $1 \leq x \leq a_i(t)a_j(t)$ because the value of x cannot be less than zero due to the max function in A_i and A_j , and it cannot exceed $a_i(t)a_j(t)$ since the counters cannot drop by more than what they started with. The latter reason also explains the bound in the inequality.

Furthermore,

$$\begin{aligned}
\Pr(1 \leq A_i^- \cdot A_j^- \leq a_i(t)a_j(t)) &\leq \Pr(A_i^- \cdot A_j^- \geq 1) \\
&\leq \Pr(\text{Both } i \text{ and } j \text{ succeed}) \\
&\leq \frac{4K}{a_i(t) + 4} \frac{4K}{a_j(t) + 4}.
\end{aligned}$$

The second inequality is due to the fact that the term $A_i^- \cdot A_j^-$ can only be positive when both stations send successfully. In order for both stations to succeed, they must at least attempt to send (and they have at most K chances to send), which is the reason for the last inequality.

Returning back to the task of bounding the term $E[A_i^- \cdot A_j^-]$, by substituting we obtain the following,

$$\begin{aligned}
E[A_i^- \cdot A_j^-] &\leq a_i(t)a_j(t) \cdot \frac{4}{a_i(t) + 4} \cdot \frac{4K^2}{a_j(t) + 4} \\
&\leq 16K^2.
\end{aligned}$$

Next we bound the value of $E[(A_i^-)^2]$ as follows.

$$\begin{aligned}
E[(A_i^-)^2] &\leq \sum_{x:1 \leq x \leq a_i(t)} x \cdot \Pr((A_i^-)^2 = x) \\
&\leq \sum_{x:1 \leq x \leq a_i(t)} (a_i(t))^2 \cdot \Pr((A_i^-)^2 = x) \\
&= (a_i(t))^2 \cdot \sum_{x:1 \leq x \leq a_i(t)} \Pr((A_i^-)^2 = x) \\
&\leq (a_i(t))^2 \cdot \Pr((A_i^-)^2 \geq 1) \\
&\leq (a_i(t))^2 \cdot \Pr(\text{User } i \text{ succeeds}) \\
&\leq (a_i(t))^2 \cdot \frac{4k(s)}{a_i(t) + 4} \\
&\leq 4a_i(t)k(s).
\end{aligned}$$

Again, we only consider values of x where the drop is positive and is at most $a_i(t)$ due to the reasons outlined above.

Substituting into Equation (5.4), we get

$$\begin{aligned}
E[(A^-)^2] &\leq \sum_{i=1}^n 4a_i(t)k(s) + 2n^2 \cdot 16K^2 \\
&< 4k(s)f(s) + 32n^2K^2.
\end{aligned}$$

Finally, we substitute these bounds into Equation (5.1) to get

$$\begin{aligned}
E[(f(s') - f(s))^2] &= (\alpha n^{3/2} \cdot nK)^2 + (nK)^2 + (\alpha n^{3/2}K)^2 + E[(A^-)^2] \\
&\quad + 2\alpha n^{3/2}K E[A^-] + 2\alpha n^{3/2} \cdot (nK)^2 \\
&\leq (\alpha n^{3/2} \cdot nK)^2 + (nK)^2 + (\alpha n^{3/2}K)^2 + 4k(s)f(s) + 32n^2K^2 \\
&\quad + 2\alpha n^{3/2}K \cdot 4nK + 2\alpha n^{3/2} \cdot (nK)^2 \\
&= 4k(s)f(s) + K^2(\alpha^2 n^5 + n^2 + \alpha^2 n^3 + 32n^2 + 8\alpha n^{5/2} + 2\alpha n^{7/2}) \\
&\leq (2\sigma - 1)k(s)f(s),
\end{aligned}$$

where the last inequality is obtained assuming that $V_0 = 64\alpha^2 n^5 K^2$ and n is sufficiently big (i.e. we started with a potential $f(s) \geq V_0$). $\square$

The following lemma shows that Condition (V2) is also satisfied.

Lemma 5.5 *Condition (V2) is satisfied for the chain X_t representing the evolution of the age-based protocol and the potential function f .*

Proof: Let the constant d in Condition (V2) be equal to $\alpha n^{5/3}$. Condition (V2) is clearly satisfied since we cannot have a bigger increase in the potential in a single step. Note that this change can only happen when the system is empty and it receives n new arrivals at once. $\square$

Next we show that if we start from state s with potential $f(s) > V_0$ and allow the protocol to run for $k(s)$ steps, then we will get a suitable negative drift which is proportional to the number of steps (i.e. the drift is $< -\sigma k(s)$). This shows that Condition (C2) in Lemma 2.8 is satisfied.

Lemma 5.6 *Given the Markov chain X_t , Condition (C2) is satisfied for the potential function f .*

Because of its length, the proof of Lemma 5.6 is presented in the next subsection.

5.1.2 Negative drift

In this section we will show that Condition (C2) is true for arrival rates $\lambda \leq 1/(\alpha n^{3/4})$. We proceed by dividing the analysis into three cases according to the starting state. For the purposes of the analysis presented below, assume that the arrival rate is $\lambda = \frac{1}{\alpha' n^{3/4}}$ for some $\alpha' \geq \alpha$. Let $m(X_t)$ be the number of users i with $q_i(t) > 0$ and counters $a_i(t) < \beta n$, and let $m'(X_t)$ be the number of users i with $q_i(t) > 0$ and counters $a_i(t) < \lceil 16n^{1/2} \rceil$. As Lemma 2.8 states, we need only consider states in $\mathbf{X} - C$, and we divide them into the following three cases.

Case 1. $m'(X_t) = 0$ and $m(X_t) < \lceil n^{1/2} \rceil$

For every state $x = X_t$ such that $m'(x) = 0$ and $m(x) < \lceil n^{1/2} \rceil$ we define $k(x) = 1$. Furthermore, for every such state X_t we will show that

$$E[f(X_{t+1}) - f(X_t)] \leq -\epsilon,$$

where $\epsilon = \sigma$.

The proof for this case depends on a careful calculation of the expected change in the potential in a single step. The intuition here is almost the same as the one used in Case 1 in Section 3.3. It depends on the fact that since all active users have large age counters, then it is likely that most will stay silent in the single step examined here. On the other hand, if one such user sends, then it is likely to succeed since the others will probably stay silent.

For convenience we will define $m = m(X_t)$, $m' = m'(X_t)$, and ℓ will denote the number of users i such that $q_i(t) > 0$. Without loss of generality assume that these nonempty queues are the queues of users $1, \dots, \ell$. Let p_i denote the probability that user i attempts to send at step t . Notice that for the nonempty users (i.e. for $1 \leq i \leq \ell$), the probability of sending is $p_i = \frac{4}{a_i+4}$. For the users which were empty just before t , there is a probability $\lambda_i = \lambda/n$ that there will be a new arrival (which is sent immediately because the user has an empty queue). Assuming symmetric arrivals, for users $\ell+1 \leq i \leq n$, the probability of sending is $p_i = \lambda/n$.

Let

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

be the probability that all of the n users stay silent at step t . Also, let us define the sum

$$S = \sum_{i=1}^n \frac{p_i}{1 - p_i}.$$

Note that $S \times T$ is the expected number of successful attempts in a single step. Let $I_{a,i,t}$ be the 0/1 indicator random variable which is 1 iff there is an arrival at user i during step t , and let $I_{s,i,t}$ be the 0/1 indicator random variable which is 1 iff user i succeeds in sending a packet at step t .

Let σ_i denote the probability that user i does not succeed in sending at step t and γ_i denote the probability that user i sends successfully at step t . The

expected change in the potential function can be calculated as

$$\begin{aligned}
& E[f(X_{t+1}) - f(X_t)] \\
&= \alpha n^{3/2} \sum_{i=1}^n (E[I_{a,i,t}] - E[I_{s,i,t}]) + \sum_{i=1}^n (E[a_i(t+1)] - a_i(t)), \\
&= \alpha n^{3/2}(\lambda - ST) + \sum_{i=1}^{\ell} (\sigma_i - a_i(t)\gamma_i) + \sum_{i=\ell+1}^n \lambda_i(1 - \frac{T}{1-\lambda_i}), \tag{5.5}
\end{aligned}$$

$$\begin{aligned}
&= \alpha n^{3/2}(\lambda - ST) + \sum_{i=1}^{\ell} (1 - p_i \frac{T}{1-p_i} - a_i(t) \cdot p_i \frac{T}{1-p_i}) \\
&\quad + \sum_{i=\ell+1}^n \lambda_i(1 - \frac{T}{1-\lambda_i}), \tag{5.6}
\end{aligned}$$

$$\begin{aligned}
&= \alpha n^{3/2}(\lambda - ST) + \ell - \sum_{i=1}^{\ell} (a_i(t) \cdot \frac{4}{a_i(t) + 4} \cdot \frac{T}{1 - 4/(a_i(t) + 4)}) \\
&\quad + \frac{(n-\ell)\lambda}{n} - \sum_{i=1}^{\ell} p_i \frac{T}{1-p_i} - \sum_{i=\ell+1}^n \lambda_i(\frac{T}{1-\lambda_i}), \\
&= \alpha n^{3/2}\lambda - \alpha n^{3/2}ST + \ell + \frac{(n-\ell)\lambda}{n} - 4\ell T - ST, \\
&= \alpha n^{3/2}\lambda + \ell + \frac{(n-\ell)\lambda}{n} - T((\alpha n^{3/2} + 1)S + 4\ell).
\end{aligned}$$

Equation (5.5) is similar to what we did in the derivation of the expected change in potential for the backoff protocols. However, age-based protocols are a little bit different since the age counter increases whenever we do not succeed; while backoff counters stay the same when the user is silent and increase only when there is a collision. Therefore, σ_i in Equation (5.5) is the probability that user i does not succeed (whether it collided or just stayed silent).

Note that we assume that the age counter of any empty user is set to zero. Therefore, we obtain Equation (5.6) by noticing that for every user $1 \leq i \leq \ell$, the probability that it does not succeed is $\sigma_i = 1 - p_i \frac{T}{1-p_i}$, and the probability that it does is $\gamma_i = p_i \frac{T}{1-p_i}$. For the empty users, the age counter will increase by one when there is an arrival and a collision (i.e. when there is an arrival into the empty user, and not all other stations stay silent). This happens with probability $\lambda_i(1 - T/(1 - \lambda_i))$.

To simplify this further, we will next calculate a bound on S and T . Notice

that,

$$\begin{aligned}
S &= \sum_{i=1}^n \frac{p_i}{1-p_i} \\
&= \sum_{i=1}^{\ell} \left(\frac{4/(a_i(t)+4)}{1-4/(a_i(t)+4)} \right) + \frac{\lambda(n-\ell)}{n-\lambda} \\
&\geq \sum_{i=1}^m \left(\frac{4}{\beta n} \right) + \frac{\lambda(n-\ell)}{n-\lambda} \\
&= \frac{4m}{\beta n} + \frac{\lambda(n-\ell)}{n-\lambda}.
\end{aligned}$$

Similarly, notice that for T ,

$$\begin{aligned}
T &= \prod_{i=1}^n (1-p_i) \\
&\geq \left(1 - \frac{4}{16n^{1/2}+4} \right)^m \left(1 - \frac{4}{\beta n+4} \right)^{\ell-m} \left(1 - \frac{\lambda}{n} \right)^{n-\ell} \\
&\geq 1 - \frac{4m}{16n^{1/2}+4} - \frac{4(\ell-m)}{\beta n+4} - \frac{\lambda(n-\ell)}{n} \\
&\geq 1 - \frac{m}{2n^{1/2}} - \frac{4(\ell-m)}{\beta n} - \frac{\lambda(n-\ell)}{n}.
\end{aligned}$$

Combining the bounds on S and T the expected change in the potential function is

$$\begin{aligned}
E[f(X_{t+1}) - f(X_t)] + \epsilon &\leq g(m, \ell) = \alpha n^{3/2} \lambda + \ell + \frac{n-\ell}{n} \lambda - \\
&\quad \left(1 - \frac{m}{2n^{1/2}} - \frac{4(\ell-m)}{\beta n} - \frac{\lambda(n-\ell)}{n} \right) \left((\alpha n^{3/2} + 1) \left(\frac{4m}{\beta n} + \frac{n-\ell}{n-\lambda} \lambda \right) + 4\ell \right) + \epsilon.
\end{aligned}$$

Next we will show that the function $g(m, \ell)$ is negative for all values of $0 \leq m \leq \lceil n^{1/2} \rceil$, and $\ell \geq m$. To do so we will prove that $g(m, m) < 0$, $g(m, n) < 0$, and $\frac{\partial^2 g(m, \ell)}{\partial \ell^2} > 0$ to show that the function is concave up.

We will handle the case when $m = 0$ similarly except that $m = \ell = 0$ is the empty state which is excluded from our analysis because it belongs to the set C . So we will replace $g(m, m) < 0$ above with the following for $m = 0$, $g(0, 1) < 0$.

The details of the proof are merely calculations and are presented below.

1. $g(m, m) < 0$.

$$\begin{aligned}
g(m, m) \times 2n^{11/4} \beta \alpha' (n^{7/4} \alpha' - 1) = \\
8m^2 - 8m n + 8m^2 n^{3/2} \alpha - 8m n^{5/2} \alpha + 2m n \beta + 10 m^2 n \beta - 2n^2 \beta \\
- 12 m n^2 \beta + 2n^3 \beta + 2m^2 n^{5/2} \alpha \beta - 2n^{7/2} \alpha \beta - 4 m n^{7/2} \alpha \beta \\
+ 2n^{9/2} \alpha \beta - 4m^2 n^{5/4} \alpha' + 8m n^{7/4} \alpha' - 8m^2 n^{7/4} \alpha' + 8m n^{11/4} \alpha' \\
- 4m^2 n^{11/4} \alpha \alpha' + 8m n^{13/4} \alpha \alpha' - 8m^2 n^{13/4} \alpha \alpha' + 8m n^{17/4} \alpha \alpha' \\
- 5 m^2 n^{9/4} \beta \alpha' + 6m n^{11/4} \beta \alpha' - 8 m^2 n^{11/4} \beta \alpha' + m n^{13/4} \beta \alpha' \\
+ 8 m n^{15/4} \beta \alpha' - m^2 n^{15/4} \alpha \beta \alpha' + 2m n^{17/4} \alpha \beta \alpha' + m n^{19/4} \alpha \beta \alpha' \\
- 2n^{11/4} \beta \epsilon \alpha' + 4m^2 n^3 (\alpha')^2 - 8m n^{7/2} (\alpha')^2 + 4m^2 n^{9/2} \alpha (\alpha')^2 \\
- 8m n^5 \alpha (\alpha')^2 + 4 m^2 n^4 \beta (\alpha')^2 - 6m n^{9/2} \beta (\alpha')^2 + 2n^{9/2} \beta \epsilon (\alpha')^2
\end{aligned}$$

The dominant term is $-8m n^5 \alpha (\alpha')^2$ for sufficiently large n and $\alpha > \beta$.

2. $g(m, n) < 0$.

$$\begin{aligned}
g(m, n) \times n^2 \beta^2 \alpha' = \\
n^{11/4} \alpha \beta^2 - 16 m^2 \alpha' + 16 m n \alpha' - 16 m^2 n^{3/2} \alpha \alpha' + 16 m n^{5/2} \alpha \alpha' \\
+ 2m^2 \sqrt{n} \beta \alpha' - 4m n \beta \alpha' - 16 m n^2 \beta \alpha' + 16 n^3 \beta \alpha' + 2m^2 n^2 \alpha \beta \alpha' \\
- 4m n^{5/2} \alpha \beta \alpha' + 2m n^{5/2} \beta^2 \alpha' - 3n^3 \beta^2 \alpha' + n^2 \beta^2 \epsilon \alpha'
\end{aligned}$$

Given that $\beta = 17$ and $\epsilon = \sigma < 3$, the dominant term here is $-3n^3 \beta^2 \alpha'$.

Note that when m approaches its upper bound for this case which is $\lceil n^{1/2} \rceil$,

then there will be other terms which are larger. However, all of these terms

will be dominated by another negative term, namely $-4m n^{5/2} \alpha \beta \alpha'$.

3. $\frac{\partial^2 g(m, \ell)}{\partial \ell^2} > 0$.

$$2 \left(4\alpha' n^{3/4} - \beta \right) \left(4\alpha' n^{7/4} - \alpha n^{3/2} - 5 \right) / \left(\alpha' \beta n^{7/4} \left(\alpha' n^{7/4} - 1 \right) \right) > 0$$

1' $g(0, 1) < 0$.

$$\begin{aligned}
g(0,1) \times (n^{7/4} \beta \alpha' (n^{7/4} \alpha' - 1) = \\
6 \beta - 7 n \beta + n^2 \beta + n^{3/2} \alpha \beta - 3 n^{5/2} \alpha \beta + n^{7/2} \alpha \beta - 20 n^{3/4} \alpha' \\
+ 4 n^{7/4} \alpha' - 4 n^{9/4} \alpha \alpha' + 4 n^{13/4} \alpha \alpha' - n^{7/4} \beta \alpha' + 4 n^{11/4} \beta \alpha' \\
+ n^{13/4} \alpha \beta \alpha' - n^{7/4} \beta \epsilon \alpha' + 16 n^{5/2} (\alpha')^2 - 3n^{7/2} \beta (\alpha')^2 + n^{7/2} \beta \epsilon (\alpha')^2.
\end{aligned}$$

The dominating term is $-3n^{7/2} \beta (\alpha')^2$, which dominates $+n^{7/2} \beta \epsilon (\alpha')^2$ since $\epsilon < 3$ and $\alpha < (3 - \epsilon)(\alpha')^2$. All other terms are asymptotically smaller.

Case 2. $m(X_t) \geq \lceil n^{1/2} \rceil$

For convenience, let us again define $m = m(X_t)$ and for this case let $k = k(X_t) = m + 2n$. In this case we will allow the system to run for k steps, and we will show that, given that we start in state X_t with $m(X_t) \geq \lceil n^{1/2} \rceil$, then

$$E[f(X_{t+k}) - f(X_t)] \leq -\sigma k.$$

The intuition is similar to the one used for Case 2 in the proof of positive recurrence of c -ary exponential backoff. Again, we are faced with the technical difficulty regarding the independence between some of the events and for this reason, it is helpful to identify “preamble steps” (steps in τ_0), “exceptional steps” (steps in τ_1), and also “following steps”. The definitions of these subsets of τ are almost identical to what was done in the proof for c -ary exponential backoff, and the reader is referred to the discussion and illustrative example presented there. However, we will basically show that outside of these excluded steps, there are many occasions on which the active users will have a good chance of succeeding. This will produce the required drop in the expected change of the potential.

Define $W = n^{1/2}/(e^{13}\beta)$, $A = \frac{6}{\alpha}n^{1/4}$, and let τ be the set of steps $\{t, \dots, t+k\}$. Let S denote the random variable equal to the number of successfully sent messages in τ , and let $p = \Pr(S \geq W)$ be the probability that during τ the

system sends at least W messages successfully. Then we have

$$\begin{aligned}
E[f(t+k) - f(t)] &\leq \alpha n^{3/2} \lambda k - \alpha n^{3/2} E[S] + \sum_{i=1}^n \sum_{t'=t+1}^{t+k} E[a_i(t') - a_i(t'-1)] \\
&\leq \alpha n^{3/2} \lambda k - \alpha n^{3/2} Wp + kn \\
&\leq \frac{3\alpha}{\alpha'} n^{7/4} - \alpha n^{3/2} Wp + 3n^2 \\
&\leq -\sigma k,
\end{aligned}$$

where the final inequality holds as long as $Wp > \frac{4}{\alpha} n^{1/2}$ and n is sufficiently large. This is true for any $p > 4 \cdot e^{13} \beta / \alpha$, which we will show next with plenty to spare.

To calculate the lower bound on p , let us define $B = \lceil W \rceil + \lceil A \rceil$, and the period $\tau_0 = t, \dots, t+m-1$. Let $\tau'(i)$ be the set of all $t' \in \tau$ such that $a_i(t') = 0$ and either

- (1) $q_i(t') > 0$ or
- (2) there is an arrival at user i at t' .

In other words $\tau'(i)$ contains all of the steps in τ where user i will attempt to send with probability 1. Let $\tau''(i)$ be the set of steps $t' \in \tau$ such that $a_i(t') = 0$ and $q_i(t'-1) > 0$.

Let τ_2 be the set of all $t' \in \tau$ such that

$$|\{(t'', i) \mid t'' \in \tau'(i) \text{ and } t'' < t'\} \cup \{(t'', i) \mid t'' \in \tau''(i) \text{ and } t'' < t'\}| \geq B.$$

Finally, let τ_1 denote the set of all $t' \in \tau - \tau_0 - \tau_2$ such that, for some i , $\tau'(i) \cap [t' - B + 1, t'] \neq \emptyset$. Let us now define $E1$ to be the event that there are at most A arrivals during τ .

Lemma 5.7 $\Pr(\overline{E1}) \leq e^{-\lambda k/3}$.

Proof: The expected number of arrivals in τ is $\lambda k \leq \lambda 3n \leq \frac{3}{\alpha'} n^{1/4}$. Notice that $A = \frac{6}{\alpha} n^{1/4} \geq 2\lambda k$. This is double the expected number of arrivals in τ . Using a Chernoff bound, the probability that there are at least $2\lambda k$ arrivals is at most $e^{-\lambda k/3}$. $\square$

We will now show that $\Pr(\mathcal{S} < W \mid E1) \leq 10^{-5}$ for sufficiently large n . We begin with the following lemma.

Lemma 5.8 *Given any fixed sequence of states $X_t, \dots, X_{t+z}$ that satisfies $t+z \in \tau - \tau_0 - \tau_1 - \tau_2$, $q_i(t+z) > 0$, and $a_i(t+z) \leq (\beta+3)n$, the probability that user i succeeds at step $t+z$ is at least $1/(e^{10}\beta n)$.*

Proof: If we assume the conditions of the lemma, then the following is true at $t+z$.

- There are no users j with $a_j(t+z) < B$. This is true because we do not consider the steps in τ_0 or τ_1 .
- There are at most B users j with $a_j(t+z) < m$. This condition is satisfied since there are m users which started with $q_j(t) > 0$ and $a_j(t) < m$, and after τ_0 all active users which did not succeed in sending or were empty and received a new arrival will have their age counters $\geq m$. Remember that at most B users succeed or have new arrivals.
- There are at most $\min\{m+B, n\}$ users j with $a_j(t+z) < \beta n$.

Therefore, the probability that user i succeeds during step $t+z$ is at least

$$\begin{aligned}
& \frac{4}{(\beta+3)n+4} \left(1 - \frac{4}{B+4}\right)^B \left(1 - \frac{4}{m+4}\right)^m \left(1 - \frac{4}{\beta n+4}\right)^{n-\min\{m+B, n\}} \\
& \geq \frac{4}{(\beta+3)n+4} \cdot \frac{1}{e^4} \frac{1}{e^4} \left(1 - \frac{4(n-\min\{m+B, n\})}{\beta n+4}\right) \\
& \geq \frac{1}{e^9((\beta+3)n+4)} \\
& \geq \frac{1}{e^{10}\beta n}.
\end{aligned}$$

□

Corollary 5.9 *Given any fixed sequence of states $X_t, \dots, X_{t+z}$ which does not violate $E1$, and satisfies $t+z \in \tau - \tau_0 - \tau_1 - \tau_2$ and where there were fewer than B successes in $t, \dots, t+z-1$, then the probability that some user succeeds at step $t+z$ is at least $\frac{1}{e^{11}\beta n^{1/2}}$.*

Proof: Since there were fewer than B successes, there are at least $m - B \geq n^{1/2}/2$ users with a nonempty queue and age counter $a_i \leq (\beta + 3)n$. For any two users i and i' , the event that user i succeeds at step $t + z$ is disjoint with the event that user i' succeeds at the same step. Note that $B \leq \lceil \frac{n^{1/2}}{e^{12}\beta} \rceil$ and $m \geq \lceil n^{1/2} \rceil$, and therefore, $m - B \geq n^{1/2}/2$. Summing for $m - B$ users, the probability that one of them succeeds at step $t + z$ is at least $\frac{m-B}{e^{10}\beta n} \geq \frac{1}{e^{11}\beta n^{1/2}}$. $\square$

Lemma 5.10 *If n is sufficiently large then $\Pr(S < W \mid E1) \leq 10^{-5}$.*

Proof: If $E1$ is true then τ_2 does not start until there has been W successes. Note that $|\tau - \tau_0 - \tau_1| \geq n$. Using Corollary 5.9 we can show that the expected number of successes in at least n independent Bernoulli trials with probability of success at least $\frac{1}{e^{11}\beta n^{1/2}}$ is $\geq \frac{n^{1/2}}{e^{11}\beta}$. For sufficiently large n , this is at least twice our goal W . A Chernoff bound shows that the probability of having fewer than W successes is at most $\exp(-\frac{n^{1/2}}{3e^{11}\beta}) \leq 10^{-5}$ for sufficiently large n . $\square$

Finally, we conclude that

$$\begin{aligned} p &\geq 1 - \Pr(\overline{E1}) - \Pr(S < W \mid E1) \\ &\geq 1 - e^{-\lambda k/3} - 10^{-5} \\ &\geq 0.9999. \end{aligned}$$

Case 3. $m'(X_t) \geq 1$ and $m(X_t) < \lceil n^{1/2} \rceil$

For every state X_t such that $m'(X_t) \geq 1$ and $m(X_t) \leq \lceil n^{1/2} \rceil$, we define $k = k(X_t) = 2\lceil n^{1/2} \rceil$. Again, let $m = m(X_t)$ and $m' = m'(X_t)$. Let $\tau = \{t, \dots, t+k-1\}$, $\tau_0 = \{t, \dots, t + \lceil n^{1/2} \rceil - 1\}$, and let S denote the number of successfully sent messages in τ . Finally, let $p = \Pr(S \geq 1)$.

Since in this case $m' \geq 1$, there is at least one user j such that $a_j(t) \leq \lceil 16n^{1/2} \rceil$ and $q_j(t) > 0$. Let $E1$ be the event that there are no new arrivals during τ .

Lemma 5.11 $\Pr(\overline{E1}) \leq 10^{-5}$.

Proof: The expected number of arrivals in k steps is $o(1)$, and therefore,

$$\Pr(E1) \geq \left(1 - \frac{\lambda}{n}\right)^{nk} \geq 1 - \lambda k \geq 1 - 10^{-5}.$$

□

The next lemma calculates the probability that user j will send successfully after τ_0 .

Lemma 5.12 *Consider the fixed sequence of states $X_t, \dots, X_{t+z}$ which does not violate $E1$ such that there are no successes during the period $[t, \dots, t+z-1]$ and $t+z \in \tau - \tau_0$. Then the probability that user j succeeds at step $t+z$ is at least $1/(18e^4 \lceil n^{1/2} \rceil)$.*

Proof: The conditions of the lemma and the protocol imply that there were no successes during $[t, \dots, t+z-1]$ and that $t+z \notin \tau_0$. Also, since we are assuming event $E1$, there are no new arrivals (which potentially may create users with very small counters). This leads us to conclude that all age counters of nonempty users are at least $\lceil n^{1/2} \rceil$, and the following is true:

- There are no users i with $a_i(t+z) < \lceil n^{1/2} \rceil$.
- There are at most $m < \lceil n^{1/2} \rceil$ users with $a_i(t+z) < \beta n$.
- For the special user number j , $q_j(t+z) > 0$ and $a_j(t+z) > \lceil 18n^{1/2} \rceil$.

Therefore, probability that user j succeeds at step $t+z$ is at least

$$\begin{aligned} & \frac{4}{\lceil 18n^{1/2} \rceil + 4} \left(1 - \frac{4}{\lceil n^{1/2} \rceil + 4}\right)^{\lceil n^{1/2} \rceil} \left(1 - \frac{1}{\beta n}\right)^n \\ & \geq \frac{4}{\lceil 18n^{1/2} \rceil + 4} \cdot \frac{1}{e^4} \left(1 - \frac{1}{\beta}\right) \\ & \geq \frac{1}{18e^4 \lceil n^{1/2} \rceil}. \end{aligned}$$

□

Lemma 5.13 *If n is sufficiently large then $\Pr(S < 1 \mid E1) \leq e^{-1/(18e^4)}$.*

Proof: There are at least $|\tau - \tau_0| \geq \lceil n^{1/2} \rceil$ steps where Lemma 5.12 shows that the probability of having no success in each step is a Bernoulli trial with probability of success $\geq 1/(18e^4 \lceil n^{1/2} \rceil)$. Therefore, the probability of having no success at all in the $|\tau - \tau_0|$ steps is at most

$$\left(1 - \frac{1}{18e^4 \lceil n^{1/2} \rceil}\right)^{\lceil n^{1/2} \rceil} \leq e^{-1/(18e^4)}.$$

□

To conclude Case 3, observe that $p \geq 1 - \Pr(\overline{E1}) - \Pr(\mathcal{S} < 1 \mid E1)$. Using Lemmas 5.11 and 5.13, this is at least $1 - 10^{-5} - e^{-1/(18e^4)} \geq 0.001$. Therefore, the expected change in the potential in k steps is

$$\begin{aligned} E[f(X_{t+k}) - f(X_t)] &\leq \alpha n^{3/2}(\lambda k - Wp) + nk \\ &\leq \alpha n^{3/2}\left(\frac{4}{\alpha' n^{1/4}} - 0.001\right) + 2n^{3/2} \\ &\leq -\sigma (2\lceil n^{1/2} \rceil) \\ &= -\sigma k. \end{aligned}$$

This concludes Case 3. Finally, we prove the main result of this section.

Proof of Theorem 5.1: Suppose that $\lambda \leq 1/(\alpha n^{3/4})$, then Lemmas 5.2, 5.6, and 5.5 show that the conditions of Lemma 2.8 are all satisfied by the Markov chain X_t and the associated potential function f . Notice also that the load of the system is $L(x) \leq f(x)$ for any state x . Thus, using Lemma 2.8 we conclude that $E[T_{ret}] < \infty$ and $E_\pi[L(X_t)] < \infty$. □

5.2 The Stability of Sublinear Age-based Protocols

In the previous section, we examined an age-based protocol where active users with age a_i send with probability $\Theta(1/a_i)$. In this section, we will consider protocols where the send probability is $\Theta(a_i^{-\varepsilon})$ for some $0 < \varepsilon < 1$. We will show that such protocols are “stable”, in the sense that the expected time to return to the empty state is finite (i.e. the chain is positive recurrent), as long as the

arrival rate is $\lambda < 1/(\alpha n^{1+1/\varepsilon})$ for some sufficiently large constant α and big enough n . Note that we use the symbol ε for the powers of the age counters and in the arrival rate, which should not be confused with the symbol ϵ used in Theorem 2.2.

The formal result regarding the stability of these age-based protocols is stated in the following theorem.

Theorem 5.14 *Suppose that a multiple-access channel is shared by n users running an age-based protocol and that user i with a nonempty queue sends with probability $p_i = (a_i^\varepsilon + 1)^{-1}$ for some constant $0 < \varepsilon < 1$. Then there exists a constant α such that if n is sufficiently large, and the overall arrival rate to the system is $\lambda \leq \frac{1}{\alpha n^{1+1/\varepsilon}}$ then the expected return time to the empty state is finite.*

Proof:

Let X_t be the Markov chain representing the system. Each state in X_t is a $2n$ -tuple holding the information about the n queue sizes and the n age counters. The potential function is defined as

$$f(X_t) = 8n^{1+1/\varepsilon} \sum_{i=1}^n q_i(t) + \sum_{i=1}^n a_i(t).$$

We will invoke Theorem 2.2 to prove that the expected waiting time is finite. Note that the Markov chain X_t satisfies the initial conditions of the theorem. That is, it is time-homogeneous, irreducible, and aperiodic and has a countable state space.

The proof will be divided into two cases depending on the magnitude of the age counters. The cases and their analysis will be similar to what we have been doing for the other protocols discussed earlier. The intuition behind the proof is as follows. In the first case we will assume that there is a nonempty queue with a small age counter. Since each active user sends with probability $(a_i^\varepsilon + 1)^{-1}$ and there might be many users with small counters, we will allow the system to run for $n^{1/\varepsilon}$ steps to raise these counters to a suitable level. Then in the next $n^{1/\varepsilon}$ steps at least one queue will have a good chance of sending

successfully. If we assume that the arrival rate is low enough, the queue part of the potential will be negative and will dominate the drift and make it also negative. The second case will depend on the large age counters and the fact that if one of them succeeds then the drop in the potential will be very large. In this case, the age counter part of the potential will dominate the other terms and the drift will become negative.

Case 1. There exists a user ξ with $q_\xi(t) > 0$ and $a_\xi(t) < n^{1/\varepsilon}$.

For every state s , such that there exists a user u with $q_u > 0$ and $a_u < \lceil n^{1/\varepsilon} \rceil$, we will define an integer k (which depends upon s) and we will show that, if $X_t = s$, then $E[f(X_{t+k}) - f(X_t)] \leq -\epsilon k$, for some constant $\epsilon > 0$.

For this case, let $k = 2\lceil n^{1/\varepsilon} \rceil$, let τ be the set of all steps $\{t, \dots, t+k-1\}$, and let S be the random variable denoting the number of successes that the system has during τ . Let $p = \Pr(S \geq 1)$. Then we have

$$\begin{aligned}
E[f(X_{t+k}) - f(X_t)] &\leq 8n^{1+1/\varepsilon}(\lambda k - E[S]) + \sum_{i=1}^n \sum_{t'=t+1}^{t+k} E[a_i(t') - a_i(t'-1)] \\
&\leq 8n^{1+1/\varepsilon}(\lambda k - p) + kn \\
&\leq 8n^{1+1/\varepsilon} \frac{1}{\alpha n^{1+1/\varepsilon}} 2\lceil n^{1/\varepsilon} \rceil - 8n^{1+1/\varepsilon} p + 2\lceil n^{1/\varepsilon} \rceil n \\
&\leq \frac{16}{\alpha} \lceil n^{1/\varepsilon} \rceil - 8n^{1+1/\varepsilon} p + 4\lceil n^{1+1/\varepsilon} \rceil \\
&\leq -\epsilon k,
\end{aligned}$$

where the final inequality holds when $p \geq 1/2$ and n and α are sufficiently big.

Thus we must show that $p \geq 1/2$. Let $E1$ be the event that there are no new arrivals into the system throughout τ , and note that

$$\Pr(E1) \geq 1 - \lambda k \geq 1 - \frac{4}{\alpha n} \geq 1 - 10^{-5}.$$

Let τ_0 denote the steps $\{t+1, \dots, t+\lceil n^{1/\varepsilon} \rceil - 1\}$.

Consider any fixed sequence of states $X_t, \dots, X_{t+z}$ where:

1. the steps do not violate $E1$,
2. $t+z \in \tau - \tau_0$,

3. $q_u(t+z) > 0$, and

4. $a_u(t+z) \leq 3n^{1/\varepsilon}$.

Then the probability that user u succeeds at step $t+z$ is

$$\begin{aligned}
p_u \cdot \prod_{j \neq u} (1 - p_j) &\geq \frac{1}{a_u^\varepsilon + 1} \prod_{j \neq u \wedge q_j \geq 1} \left(1 - \frac{1}{a_j^\varepsilon + 1}\right) \prod_{j \neq u \wedge q_j = 0} (1 - \lambda_j) \\
&\geq \frac{1}{3^\varepsilon n + 1} \prod_{j=1}^{n-1} \left(1 - \frac{1}{n}\right) \\
&\geq \frac{1}{4n + 1} \cdot \frac{1}{4} \\
&\geq \frac{1}{32n}.
\end{aligned} \tag{5.7}$$

Inequality (5.7) is due to the fact that for all empty users j , the probability of getting an arrival is $\lambda_j \leq 1/n$. Also, all nonempty queues other than queue u will have their age counters reach at least $\lceil n^{1/\varepsilon} \rceil$ since $t+z \notin \tau_0$. Finally, there are $n-1$ users other than user u .

Since there are at least $|\tau - \tau_0| = \lceil n^{1/\varepsilon} \rceil$ steps satisfying Conditions 1-4, $\Pr(S \geq 1)$ is bounded from below by the probability of success in $n^{1/\varepsilon}$ Bernoulli trials with probability of success $\frac{1}{32n}$ in each trial. Thus, for n sufficiently big,

$$\begin{aligned}
\Pr(S < 1 \mid E1) &\leq \left(1 - \frac{1}{32n}\right)^{\lceil n^{1/\varepsilon} \rceil} \\
&\leq e^{-(n^{1/\varepsilon-1})/32} \\
&\leq 10^{-5}.
\end{aligned}$$

This shows that $p \geq 1 - \Pr(S < 1 \mid E1) - \Pr(\overline{E1}) \geq 1/2$ and we are done with this case.

Case 2. For all users i , if $q_i(t) > 0$ then $a_i(t) \geq n^{1/\varepsilon}$

A single step will be sufficient to show that the drift is negative in this case. Thus for all states s in this case, $k(s) = 1$. Given any state s where all users i with $q_i > 0$ have $a_i \geq n^{1/\varepsilon}$, we will show that

$$E[f(X_{t+1}) - f(X_t)] \leq -\epsilon.$$

Let ℓ denote the number of users i with $q_i \geq 1$, and assume without loss of generality that these nonempty users are users $1, \dots, \ell$. In a single step, the total queues size is expected to increase by λ , and this gets multiplied by $8n^{1+1/\varepsilon}$ when calculating the expected change in the potential. Meanwhile, the age counters will increase by

$$\begin{aligned} \sum_{i=1}^n a_i(t+1) - a_i(t) &\leq \sum_{i=1}^{\ell} 1 + \sum_{i=\ell+1}^n \lambda_i \\ &\leq \ell + \lambda. \end{aligned}$$

The first inequality is justified by the fact that the age of nonempty users can increase by at most 1, while the empty user's age increases by 1 when it receives an arrival.

Since there are ℓ users with $q_i > 0$ and $a_i \geq n^{1/\varepsilon}$, each of user i has an age counter which is expected to decrease (when it succeeds) by at least

$$\begin{aligned} &p_i \prod_{j \neq i} (1 - p_j) (a_i) \\ &\geq \frac{1}{a_i^\varepsilon + 1} \left(1 - \frac{1}{n}\right)^{n-1} (a_i) \\ &\geq \frac{a_i}{2a_i^\varepsilon} \cdot \frac{1}{4} \\ &= \frac{1}{8} \cdot a_i^{1-\varepsilon} \\ &\geq \frac{1}{8} \cdot n^{1/\varepsilon-1}. \end{aligned}$$

Therefore, the expected change in the potential in a single step is

$$\begin{aligned} E[f(X_{t+1}) - f(X_t)] &\leq 8n^{1+1/\varepsilon}\lambda + \ell + \lambda - \sum_{i=1}^{\ell} \Pr(\text{User } i \text{ succeeds}) \cdot (a_i) \\ &\leq \frac{8}{\alpha} + \ell + \lambda - \ell \cdot \frac{1}{8} \cdot n^{1/\varepsilon-1} \\ &\leq -\epsilon \end{aligned}$$

where the last inequality is true because $0 < \epsilon < 1$ and n is sufficiently large.

This concludes Case 2 of the proof.

Now we are ready to show that the Markov chain X_t is positive recurrent. Let us define the set C in Theorem 2.2 as the set consisting of the empty state

only. The function k is defined as $k(s) = 2\lceil n^{1/\varepsilon} \rceil$ for states s in Case 1, and $k(s) = 1$ for states in Case 2. By the analysis in the previous two cases, we have shown that Condition (2.4) is satisfied. Clearly, Condition (2.5) is satisfied by the set C and the potential function f . Thus we conclude using Theorem 2.2 that the Markov chain is positive recurrent, and thus the expected return time to the start state is finite. $\square$

Chapter 6

Remarks and Open Problems

The review of related results that was presented in Chapter 1 showed that most of the efficient and well studied protocols for the multiple-access channel problem require ternary feedback and at least some limited sensing to achieve high arrival rates. On the other hand, other results indicate that some protocols perform well when they obtain some information about the number of users in the system. The only exception to these results (in the sense that the protocol was purely acknowledgement-based) was the polynomial backoff protocol analysed by Håstad, Leighton and Rogoff. This result and the excellent paper it appeared within caused a revival of interest into acknowledgement-based protocols. This was coupled with some new applications for these protocols such as optical parallel computing [24, 37] which motivated more research [22].

Furthermore, we presented in Chapter 1 some practical situations where it was necessary to use an acknowledgement-based protocol. This was evident in examples such as TCP/IP networks and web applications where the client had no way of monitoring the conflicts at the server when it did not participate in these conflicts. Such full-sensing would undoubtedly overwhelm the resources of a distributed network such as the Internet. Thus, acknowledgement-based protocols are the only solution in applications which have one or more of the following properties:

1. the population size is unknown and may vary immensely over time,
2. only acknowledgements are returned by the channel or server, and
3. users may join or leave the contention process and thus cannot use full sensing which requires knowledge of the full history.

Unfortunately, there are many open questions regarding the stability and performance of acknowledgement-based protocols. This is true even for popular protocols such as binary exponential backoff, and especially in the (more realistic) finite model. The results of Goodman, Greenberg, Madras, and March [25] were a starting point that obtained quite a low value for the stable arrival rate of binary exponential backoff. Other results were obtained later by Håstad, Leighton, and Rogoff [27] regarding the stability of polynomial backoff protocols. However, the only provable bound on the expected average load for this protocol is exponential. Also, the gap between the upper and lower bounds is still large. We have tried in this thesis to decrease this gap, with some substantial success for the upper bound on the stability rate. However, the gap is still large between $1/2$ and $\Theta(n^{3/4+\epsilon})$.

Thus, perhaps the most important open question posed at the end of this study about this class of protocols is the following: Does there exist a purely acknowledgement-based protocol which is stable for a *constant* arrival rate and which has small (say polynomial in n) expected load or waiting time?

6.1 Remarks and Open Problems on the Protocols

In this section we present some remarks on the methods used in this thesis for studying contention resolution protocols. We also include some open questions which we could not answer or which were created by our investigation of these protocols.

6.1.1 Exponential Backoff Protocols

The results in Chapters 2 and 3 are the core of this study, and their development has taken most of the time spent working toward this thesis. The proof in Chapter 3 which was originally for binary exponential backoff only, grew and grew in a manner resembling the giant component. It started as a proof of positive recurrence of binary exponential backoff for a much lower arrival rate. Then improvements in λ were made slowly, until it resulted in the paper by Al-Ammal, Goldberg, and MacKenzie [3]. Then it developed to a proof for the general c -ary class for any $c > 1$. However, even though the improvement in the stability rate is substantial over the bound discovered by Goodman et al. [25], we still did not answer the question whether there exists a constant λ^* such that the protocol is stable for $\lambda \leq \lambda^*$. This, in my opinion, is a very important open problem, especially since the best known upper bound for stability is $1/2$. It would also be nice if the form of stability of binary exponential backoff was strengthened from positive recurrence to strong stability.

We have discovered that in general, there are two main difficulties with proving stability for backoff protocols using natural potential functions similar to the ones used in this thesis.

1. Because of the backoff counters part of the potential, the single step jumps in the potential are unbounded. This is unfortunately true for both negative and positive jumps, and this eliminates the possibility of using many useful theorems for showing either stability or instability which require some bound on the jumps (such as the results in [15]).
2. The backoff counters change stochastically. Hence, after a known number of steps, their values are not exactly known. This was handled by obtaining large deviation probabilistic bounds on the counters in Lemma 3.2. These difficulties are also present on the instability side, where again bounded jumps are necessary for showing transience.

The analysis that we presented for c -ary exponential backoff (including

the case divisions and the given scenario in each case) clarifies some of the mechanisms responsible for the protocol's stability. For example, in Case 2 of the proof for c -ary backoff (which is illustrated by an example in Figure 3.1), we noticed that every time a large number of users have small backoff counters, they will collide until they back off by increasing their counter's value (this is what happens in τ_0). Then there will be a period of time where most of these users will send successfully if they run for a reasonable amount of time (steps in $\tau - \tau_0 - \tau_1 - \tau_2$). This period may be briefly interrupted by a message succeeding or a message arriving into an empty user. However, although this will create a user with a very small backoff counter, it is not in itself a "bad" thing. In our analysis we ignore such users (and the steps after they are created) until they back off from sending with such high probability (these steps are excluded by placing them in τ_1).

To be able to show that the protocol is stable for a constant arrival rate (if that is in fact true), perhaps it would be wiser to consider carefully what happens in τ_1 . When a new arrival is received by an empty user, all other active users have counters which are sufficiently high for this new message to be sent immediately with probability close to $1/e$. Also, when some user succeeds, it does not necessarily create conflict, since it has a high probability of sending the rest of its messages successfully while other users are backing off from sending. This creates an effect similar to the capture effect utilised by Håstad, Leighton, and Rogoff [27] for showing that polynomial backoff is stable. However, in our analysis, we completely ignore this effect and choose the worst possible scenario. Although this simplifies the analysis, it also (perhaps) lowers the stability arrival rate. We should also note that the capture effect is noticeable in the exploratory simulations that were executed for the exponential backoff protocol.

The reader might suggest doing some experimental work toward investigating this problem. However, simulations can be quite misleading because of the following reasons:

1. Simulations for polynomial backoff presented in Table I of [27] show a huge amount of average load after 10 million steps for arrival rates bigger than $1/2$. However, the same paper presents a proof showing that polynomial backoff is stable for any $\lambda < 1$.
2. Aldous states that simulations and heuristic arguments suggest that for many backoff functions, the protocol enters a finite but very long period where the load is very low, before instability sets in and the system load explodes [4]. When this is the case, even very long simulations might not be sufficient for showing that the system load increases unboundedly.

Thus, simulations cannot give an accurate picture of a protocol's stability. However, such experiments can help guide the researcher toward gaining some intuition which may be used in a formal proof.

Regarding the relationship between the finite and infinite models, we notice that the queues in the finite model seem to have a stabilising effect. Intuitively, they stabilise the system by limiting the number of stations in contention to a maximum of n . The binary exponential backoff protocol is known to be unstable in the infinite model for any arrival rate. Kelly and Kelly and MacPhee [34, 35] showed this for $\lambda > \ln 2$, and Aldous [4] showed that it holds for all positive λ . However, note that it can be misleading to view the infinite model as the limit (as n tends to infinity) of the finite model. For example, the polynomial backoff protocol is known to be unstable (for any positive λ) in the infinite model [34, 35], but it is stable (for any $\lambda < 1$) in the finite model [27]. Thus, Aldous's result regarding the instability of binary exponential backoff does not rule out the possibility that there is a positive constant λ^* such that binary exponential backoff protocol in the finite model is stable whenever $\lambda < \lambda^*$.

On the instability of c -ary exponential backoff, Theorem 4.2 does not do better than the result obtained by Håstad, Leighton, and Rogoff [27] for the binary case. This is mainly due to the fact that both results depend on showing that *all* states with large enough potential have nonnegative drift. Consequently, the weakest such state is the state when there exists a nonempty user with backoff

counter equal to 0, and all other users are empty. Thus, with relatively high probability, this active user will succeed in sending, and the drift will be negative. This situation required the arrival rate to be greater than λ_0 (where λ_0 is the solution for $\lambda_0 = e^{-\lambda_0}$) to ensure nonnegative drift. It remains to be seen whether a multiple step analysis can bypass this and similar cases. Certainly, Håstad et al. [27] used such an analysis with Aldous's result for the infinite case to reduce λ to at least $1/2$ for instability.

Finally, when a number of protocols are stable for a sufficiently large arrival rate, measuring the expected load or waiting time for these protocols becomes vital. Although Håstad, Leighton, and Rogoff [27] showed that polynomial backoff is stable for any $\lambda < 1$, they claim that their analysis can be used to calculate the upper bound $E[L_{avg}] = O(2^{Q(n)})$, where $Q(n)$ is a polynomial function in n . It is clear from the instability results for binary exponential backoff that it is unstable for $\lambda > 1/2$ [27]. This leads us to another open question regarding the performance of backoff (and in general acknowledgement-based) protocols under a stable arrival rate. In particular, it would be crucial to obtain polynomial or sub-polynomial bounds on the expected load of the system or the expected waiting times of messages. More on this question in the next section where we give some remarks on the age-based protocols studied in this thesis.

6.1.2 Age-based Protocols

The age-based protocol studied in Section 5.1 was similar to binary exponential backoff and the way it increased its send probabilities. These similarities enabled us to apply essentially the same methods for showing its stability. However, the potential function and the underlying Markov chain had some properties which made the analysis simpler than binary exponential backoff. The main such property was the ability to bound the upward jumps in the potential function. This enabled us to obtain a stronger form of stability than positive recurrence.

Returning back to the issue of getting performance bounds for the proto-

col when it is stable, the following open question was also posed by Håstad et al. [27]:

“It would be nice to determine the behaviour of $E[L_{avg}]$ 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 be possible to decide which polynomial backoff protocol is best (i.e., which minimised $E[L_{avg}]$ for particular λ and n .)”

It might be possible to give a partial answer this question using a technique proposed by Bertsimas, Gamarnik, and Tsitsiklis [5] which can calculate the expected value of the potential function in the steady state of the Markov chain. Their method works for any irreducible, countable state space, homogeneous and aperiodic Markov chain. However, they assume that the chain is positive recurrent with stationary distribution π and that the expected value of the potential function is bounded in the stationary distribution ($E_\pi[f(X_t)] < \infty$). Furthermore, they assume that the jumps of this potential function are bounded in the absolute value by a constant ν , and that if the potential exceeds some constant V_0 , then the drift is provably $\leq -\gamma$. If all of these conditions are met, then the expected value of the potential function is $E_\pi[f(X_t)] \leq V_0 + 2\nu^2/\gamma$.

The nice property about this method is that as a by-product of a stability analysis, the researcher usually calculates the value of γ and V_0 . Thus we can obtain bounds on the expected value of the potential function (and therefore on the expected load or the waiting time, given that the potential is an upper bound on these measures).

Unfortunately, we could not show that for exponential backoff protocols the expected value of the potential function is bounded. However, we have shown that $E_\pi[f(X_t)] < \infty$ for the linear age-based protocol defined in Section 5.1. There are two other problems with applying the methods of Bertsimas et al. to obtain a bound on the expected value of the potential. First, the jumps of our potential function are not bounded by a constant ν . This problem can be overcome by calculating a bound on the upward jumps ν_{up} and an expected

bound on the downward jumps ν_{down} . Setting $\nu = \max\{\nu_{up}, \nu_{down}\}$ will make their proof work for our Markov chain. Secondly, the drift they use (i.e. γ) is calculated for a single step, while we need multiple steps to get negative drift. This problem can be fixed by extending their proof using an embedded chain method such as the one used in Chapter 2.

Given that these changes are made to their proof, and noting that V_0 and ν are polynomial in n , and that the drift is inverse polynomial, it seems that we can show that $E_\pi[f(X_t)] < poly(n)$. Since the potential function is an upper bound on the ages of the messages in the system, this can also give us a polynomial bound on the expected waiting time in the steady state.

By studying protocols with $p_i = i^{-\epsilon}$ we were attempting to achieve better stability. Intuitively, comparing polynomial and exponential backoff protocols and their stability indicates that using a function slower than exponential would lead to more stability. Hastad et al. [27] argue that exponential backoff isn't stable for high arrival rates because it backs off too much, and that researchers should study protocols where the probability of retransmission decreases more slowly. However, the provable stable arrival rate obtained in Theorem 5.14 gets worse as ϵ approaches 0. We do not think that this reflects what is truly happening, and is probably caused by the specific potential function used in the proof. Perhaps a capture argument such as the one used in [27] will work better for the problem in Section 5.2.

6.1.3 The Search for n

A major part of the design of any contention resolution protocol is concerned with finding an approximate value of the number of users participating in the contention process. In the case of m -ary feedback, this information is immediately available and protocols can perform very well. However, m -ary feedback is usually not available for many practical situations. Thus, a protocol must have some mechanism for approximating n from the limited binary or ternary feedback available to it.

others” [42, page 25]. To solve this he proposes a limited sensing approach to the problem. Users (or hosts) do not need to monitor the channel from step 1, but when they become active they listen to collisions even when they do not attempt to send. He claims that when n users are in contention, the value of n will be reached in a logarithmic amount of time (instead of linear time in the case of binary exponential backoff).

Considering these remarks, it would be interesting to investigate the stability of such a *limited sensing* protocol, especially since this type of feedback is available in many broadcast-like channels such as local area networks. Although tree protocols are stable under the limited sensing model, each user joining the system must enter a lengthy synchronisation period. Furthermore, a backoff-like protocol with limited sensing will, in general, be much simpler than a tree protocol.

If n can be approximated in a logarithmic number of steps by such a protocol, then the length of the initialisation steps in τ_0 (in the analysis of Case 2 in Chapter 3) could be very small indeed. This in turn results in a very small number of arrivals which must be sent by the protocol to make the drift negative. Furthermore, the length of the period τ_1 will also be reduced. However, note that such a protocol is no longer acknowledgement-based, since we must assume that active users can sense other users’ collisions.

6.2 Other Techniques for Showing Stability

Our principal technique for showing that a contention resolution system is stable has been to model it using a Markov chain, and then to design a potential function that will produce negative drift. The use of Markov chains to model the system is only possible when the arrival process has a well defined distribution which adheres to the Markov property (such as a Bernoulli or a Poisson distribution). We also assumed that the arrivals to each user are independent of the other users. However, both of these restrictions are not available in many

practical situations. The arrivals may be interdependent and their distribution may not be known or may change after a certain number of steps.

Furthermore, there have been some studies that suggest that the traffic in actual Ethernet networks has a self-similar nature [36]. Such traffic is likely to display similarities across a wide range time scale. The arrivals are not just dependent, but may be so on a large scale. The question, therefore, is how do we show stability in such situations (i.e. when the underlying process is not even a Markov chain)?

If stability is defined as having a bounded expected load only (and nothing is said about the existence of a stationary distribution), then there are techniques for proving this kind of stability for non-Markovian systems. The problem of showing that a sequence of random variables $X_0, X_1, \dots$ with real valued domain is guaranteed to have $\sup_t E[X_t] < \infty$ was first studied by Hajek [29]. He showed that two conditions were necessary for this sequence of random variables to have bounded expectation. First, the sequence must have bounded exponential moments, and secondly, whenever the value of X_t is large enough it must have negative drift. When these conditions are met, then Hajek showed that $\sup_t E[X_t] < \infty$.

The results of Hajek received increasing attention after the introduction of Adversarial Queueing Theory by Borodin et al. [6]. The authors of this paper studied routing and queueing systems under worst conditions and without any assumptions of independence of arrivals or use of a fixed arrival distribution. Thus they needed a tool for showing stability in a more general context than Markov chain stability. Motivated by this research, Pemantle and Rosenthal [45] showed that such stability is possible under more relaxed condition than a bounded exponential moment. They proved the following theorem.

Theorem 6.1 (Pemantle and Rosenthal [45]) *Let X_t be a sequence of random variables and suppose that there exists constants $a > 0$, $J, V < \infty$, and $p > 2$, such that $X_0 \leq J$, and for all n ,*

$$E[X_{t+1} - X_t \mid \mathcal{F}_t] \leq -a \quad \text{on the event} \quad \{X_t > J\} \quad (6.1)$$

and

$$E[|X_{t+1} - X_t|^p | X_0, \dots, X_t] \leq V. \quad (6.2)$$

Then for any $r \in (0, p - 1)$ there is a $c = c(p, a, V, J, r) > 0$ such that $E[X_n^+] < c$ for all t .

The major advantage of these conditions is that they say nothing about the independence or memorylessness of the sequence X_t . Thus they can be applied to any random variable sequence (not necessarily Markovian). Given a contention resolution system with states Y_t , define a potential function $f(Y_t)$ which is an upper bound on the system load and set $X_t = f(Y_t)$. Clearly, if we can show that $\sup_t E[X_t] < \infty$, then the expected load will also be bounded.

From the initial investigation that we did, it seems that the negative drift condition in this theorem is satisfied by the protocols that we studied earlier for some sufficiently small mean arrival rate. Unfortunately, however, we could not apply these results to the study of backoff and age-based protocols. The main problem seems to be in finding a potential function that satisfies Condition (6.2). This condition seems very strong in the case of our potential functions. For exponential backoff, we cannot satisfy Condition (6.2) when $p > 1$. Even in age-based protocols, where positive jumps are bounded by a constant, negative jumps resulting from a huge age counter being reset to zero are unbounded. Perhaps this is another reason for using bounded jump functions (if they exist) for showing stability or instability, since a function with bounded jumps will trivially satisfy this condition.

Another method for proving stability uses fluid limit models. This method is different from the Foster-type conditions which depend on getting negative drift for the potential function. Instead of examining the original discrete Markov chain, the fluid limit model constructs a continuous process which is a scaled limit of the original process. Dai [11] showed that if we can prove that, starting from any initial condition this continuous process eventually reaches 0, then the original system is positive recurrent. This technique has been used for showing stability for Multiclass Queueing Networks [11], problems similar to

bin packing [10], and Adversarial Queueing Networks [19]. It is an interesting but still open question whether the same technique can be used to show stability for contention resolution protocols such as the ones studied here. The major obstacle toward using these methods might be the transformation of problems such as the ones studied here from the discrete world to the continuous world of the fluid limit models.

Bibliography

- [1] N. Abramson, *The ALOHA system - another alternative for computer communications*, 1970 Fall Joint Conference AFIPS **27** (1970), 281–285.
- [2] ———, *Development of the ALOHANET*, IEEE Transactions on Information Theory **IT-31** (1985), 119–123.
- [3] H. Al-Ammal, Leslie Ann Goldberg, and Phil MacKenzie, *Binary exponential back-off is stable for high arrival rates*, Proceedings of the International Symposium on Theoretical Aspects of Computing **17** (2000), 169–180, to appear in Theory of Computing Systems.
- [4] D. Aldous, *Ultimate instability of exponential back-off protocol for acknowledgement-based transmission control of random access communication channels*, IEEE Transactions on Information Theory **IT-33** (1987), no. 2, 219–223.
- [5] D. Bertsimas, D. Gamarnik, and J. N. Tsitsiklis, *Geometric bounds for stationary distributions of infinite Markov chains via Lyapunov functions*, LIDS Technical Report (August 1998), available on the web at <http://truth.mit.edu/~jnt/gamarnik.ps>.
- [6] A. Borodin, J. Kleinberg, P. Raghavan, M. Sudan, and D. P. Williamson, *Adversarial queueing theory*, Proceedings of the ACM Symposium on the Theory of Computing **28** (1996), 376–385.
- [7] P. Brémaud, *Markov chains: Gibbs fields, Monte Carlo simulation, and queues*, Springer-Verlag, New York, 1999.
- [8] J. I. Capetanakis, *The multiple access broadcast channel: protocol and capacity considerations*, Tech. Report ESL-R-806, MIT, Electronic Systems Laboratory, 1978.

- [9] A. Carleial and M. E. Hellman, *Bistable behaviour of ALOHA-type systems*, IEEE Transactions on Communications **COM-23** (1975), no. 4, 401–409.
- [10] E. G. Coffman and A. Stolyar, *Bandwidth packing*, Algorithmica (to appear).
- [11] J. G. Dai, *On positive harris recurrence of multiclass queueing networks: A unified approach via limit models*, The Annals of Applied Probability **5** (1995), no. 1, 49–77.
- [12] M. Drmota and U. Schmid, *The analysis of the expected successful operation time of slotted ALOHA*, IEEE Transactions on Information Theory **IT-39** (1993), no. 5, 1567–1577.
- [13] G. Fayolle, E. Gelenbe, and J. Labetoulle, *Stability and optimal control of the packet switching broadcast channel*, Journal of the ACM **24** (1977), no. 3, 375–386.
- [14] G. Fayolle, E. Gelenbe, J. Labetoulle, and D. Bastin, *The stability problem of broadcast packet switching*, Acta Informatica **4** (1974), 49–53.
- [15] G. Fayolle, V. A. Malyshev, and M. V. Menshikov, *Topics in the constructive theory of countable Markov chains*, Cambridge University Press, Cambridge, 1995.
- [16] R. Fielding, J. Gettys, J. C. Mogul, H. Frystyk Nielsen, and T. Berners-Lee, *Hypertext transfer protocol : http / 1.1*, W3C Journal **4** (1996), no. 1.
- [17] F. G. Foster, *On the stochastic matrices associated with certain queueing processes*, Annals of Mathematical Statistics **24** (1953), 355–360.
- [18] R. G. Gallager, *Conflict resolution in random access broadcast networks*, Proceedings of the AFOSR Workshop in Communication Theory and Applications **35** (1978), no. 8, 74–76.
- [19] D. Gamarnik, *Stability of adversarial queues via fluid models*, The 1998 IEEE Symposium on Foundations of Computing (1998), 60–70.
- [20] L. Goldberg, *Notes on contention resolution*, Available on the web at <http://www.dcs.warwick.ac.uk/~leslie>, 2000.
- [21] L. Goldberg, M. Jerrum, S. Kannan, and M. Paterson, *A bound on the capacity of backoff and acknowledgement-based protocols*, ICALP 2000 (2000).
- [22] L. Goldberg and P. MacKenzie, *Analysis of practical backoff protocols for contention resolution with multiple servers*, Journal of Computer and System Sciences **58** (1999), 232–258.

- [23] L. A. Goldberg, P. D. Mackenzie, M. Paterson, and A. Srinivasan, *Contention resolution with constant expected delay*, Journal of the ACM, to appear (2000).
- [24] Leslie Ann Goldberg, Mark Jerrum, Tom Leighton, and Satish Rao, *Doubly logarithmic communication algorithms for optical communication parallel computers*, SIAM Journal on Computing **26** (1997), no. 4, 1100–1119.
- [25] J. Goodman, A. G. Greenberg, N. Madras, and P. March, *Stability of binary exponential backoff*, Journal of the ACM **35** (1988), no. 3, 579–602.
- [26] G. R. Grimmett and D. R. Stirzaker, *Probability and random processes*, Oxford University Press, Oxford, 1992.
- [27] J. Håstad, T. Leighton, and B. Rogoff, *Analysis of practical backoff protocols for multiple access channels*, SIAM Journal on Computing **25** (1996), no. 4, 740–774.
- [28] B. Hajek, *Acknowledgement based random access transmission control: An equilibrium analysis*, Proceedings of the International Conference on Communications (1982), 1C.1.1–1C.1.7.
- [29] ———, *Hitting time and occupation-time bounds implied by drift analysis with applications*, Advances in Applied Probability **14** (1982), 502–525.
- [30] IETF, *Request for comment RFC 1122*, Available on the web at <http://www.cis.ohio-state.edu/htbin/rfc/rfc1122.html>, October 1989, Internet Engineering Task Force, Braden, R. (ed.).
- [31] M. J. Ingenoso, *Stability analysis for certain queueing systems and multi-access communication channels*, Ph.D. thesis, University of Wisconsin-Madison, 1993.
- [32] V. Jacobson, *Congestion avoidance and control*, Proceedings of SIGCOMM '88 (1988), 1.
- [33] M. Kaplan, *A sufficient condition for nonergodicity of a Markov chain*, IEEE Transactions on Information Theory **IT-25** (1979), no. 4, 470–471.
- [34] F. P. Kelly, *Stochastic models of computer communication systems*, Journal of the Royal Statistical Society B **47** (1985), 379–395.
- [35] F. P. Kelly and I. M. MacPhee, *The number of packets transmitted by collision detect random access schemes*, The Annals of Probability **15** (1987), no. 4, 1557–1568.
- [36] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson, *On the self-similar nature of Ethernet traffic*, IEEE/ACM Transactions on Networking **2** (1994), 1–15.

- [37] P. D. MacKenzie, C. G. Plaxton, and R. Rajaraman, *On contention resolution protocols and other probabilistic phenomena*, Journal of the ACM **45** (1998), no. 2, 324–378.
- [38] I. M. MacPhee, *On optimal strategies for stochastic decision processes*, Ph.D. thesis, University of Cambridge, 1987.
- [39] J. L. Massey (Ed.), *Special issue on multiple access protocols*, IEEE Transactions on Information Theory **IT-31** (1985).
- [40] R. Metcalfe and D. Boggs, *Ethernet: Distributed packet switching for local computer networks*, Communications of the ACM **19** (1976), 395–404.
- [41] S. P. Meyn and R. L. Tweedie, *Markov chains and stochastic stability*, Springer-Verlag, New York, 1993.
- [42] L. M. Molle, *A new binary logarithmic arbitration method for ethernet*, Tech. Report CSRI-298, Computer Systems Research Institute, University of Toronto, July 1994, available on the web at <ftp://ftp.cs.utoronto.ca/pub/blam/report.ps>.
- [43] L. M. Molle and George C. Polyzos, *Contention resolution algorithms and their performance analysis*, Tech. Report CS93-300, University of California at San Diego, 1993, available on the web at <http://www-cse.ucsd.edu/groups/csl/pubs/tr/CS93-300.ps>.
- [44] J. Mosely and P. A. Humblet, *A class of efficient contention resolution algorithms for multiple access*, IEEE Transactions on Communications **33** (1985), no. 2, 145–151.
- [45] R. Pemantle and J. S. Rosenthal, *Moments conditions for a sequence with negative drift to be uniformly bounded in L^r* , Stochastic Processes and Their Applications **82** (1999), 143–155.
- [46] P. Raghavan and E. Upfal, *Contention resolution with short delays*, Proceedings of the ACM Symposium on the Theory of Computing **24** (1995), 229–237.
- [47] W. A. Rosenkrantz, *Some theorems on the instability of the exponential back-off protocol*, PERFORMANCE '84 (1984), 199–205.
- [48] W. A. Rosenkrantz and D. Towsley, *On the instability of the slotted ALOHA multiaccess algorithm*, IEEE Transactions on Automatic Control **AC-28** (1983), 994–996.
- [49] F. M. Spieksma and R. L. Tweedie, *Strengthening ergodicity to geometric ergodicity for markov chains*, Stochastic Models **10** (1994).

- [50] W. R. Stevens, *TCP/IP illustrated, volume 1: The protocols.*, Addison-Wesley, Reading MA, USA, 1994.
- [51] J. Stidham, *The last word on the $L=\lambda W$* , Operations Research **22** (1974), 417–421.
- [52] R. S. Tsybakov and V. A. Mikhailov, *Free synchronous packet access in a broadcast channel with feedback*, Problemy Peredachi Informatsii **14** (1978), no. 4, 259–280.
- [53] ———, *Ergodicity of a slotted ALOHA system*, Problems of Information Transmission **15** (1979), 301–312.
- [54] ———, *Random multiple packet access: Part-and-try algorithm*, Problemy Peredachi Informatsii **16** (1980), no. 3, 80–94.