

Stochastic models for resource sharing in wireless networks

Citation for published version (APA):

Ven, van de, P. M. (2011). Stochastic models for resource sharing in wireless networks. [Phd Thesis 1 (Research TU/e / Graduation TU/e), Mathematics and Computer Science]. Technische Universiteit Eindhoven. https://doi.org/10.6100/IR719838

DOI: 10.6100/IR719838

Document status and date:

Published: 01/01/2011

Document Version:

Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

• A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.

• The final author version and the galley proof are versions of the publication after peer review.

• The final published version features the final layout of the paper including the volume, issue and page numbers.

Link to publication

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- · Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
 You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

Stochastic Models for Resource Sharing in Wireless Networks





CIP-DATA LIBRARY TECHNISCHE UNIVERSITEIT EINDHOVEN

Ven, Peter M. van de

Stochastic Models for Resource Sharing in Wireless Networks / by P.M. van de Ven – Eindhoven : Technische Universiteit Eindhoven, 2011. A catalogue record is available from the Eindhoven University of Technology Library ISBN : 978-90-386-3010-6 NUR 919 Subject headings : queueing theory, wireless networks 2000 Mathematics Subject Classification : 60K25, 68M20, 90B18

Printed by Ponsen & Looijen. Cover by Oranje Vormgevers.

Stochastic Models for Resource Sharing in Wireless Networks

PROEFSCHRIFT

ter verkrijging van de graad van doctor aan de Technische Universiteit Eindhoven, op gezag van de rector magnificus, prof.dr.ir. C.J. van Duijn, voor een commissie aangewezen door het College voor Promoties in het openbaar te verdedigen op maandag 19 december 2011 om 16.00 uur

door

Peter Maria van de Ven

geboren te Oss

Dit proefschrift is goedgekeurd door de promotor:

prof.dr.ir. S.C. Borst

Copromotor: dr. J.S.H. van Leeuwaarden

ACKNOWLEDGEMENTS

This thesis marks the end of my four-year career as a Ph.D. student, and I would like to take this opportunity to thank the many people that helped me to bring this period to a good end.

First, I would like to thank Sem Borst and Johan van Leeuwaarden for the very pleasant and instructive collaboration; it would be difficult to overstate how lucky I have been having had both as an advisor. Sem has helped me to take my first steps in doing research, and I benefitted immensively from his enviable intuition and impressive persistence. Johan's enthusiasm and energy have been infective, and I would usually emerge from our meetings feeling inspired and productive.

During the past four years I was lucky enough to work with, and learn from, a variety of people, and I would like to thank the other co-authors of the papers on which this thesis was build: Dee Denteneer, Alexandre Proutière, Seva Shneer and Lei Ying. A special word of thanks goes to Guido Janssen, who co-authored the papers resulting in Chapters 5-7 and carefully proofread this manuscript, and in general has proven to be a very instructive collaborator. I am also grateful to Nidhi Hegde, Laurent Massoulié and Theodoros Salonidis for giving me the opportunity to spent some very interesting months at Technicolor Paris Research Lab.

I am indebted to Richard Boucherie, Remco van der Hofstad, Jean-Paul Linnartz, Rob van der Mei and Sindo Núñez Queija for agreeing to serve on my doctoral committee and for reading and commenting on this thesis. Marko Boon helped me out on many occasions with IT-related issues, for example creating the template on which this thesis was built. I would also like to thank Niek Bouman and Maren Eckhoff for proofreading the introduction chapter.

In addition to being productive, my time spent as a Ph.D. student was also very pleasant, and I would like to express my gratitude to all my colleagues at Eurandom and the Stochastic OR group at TU/e for creating such a wonderful working environment. In particular, I would like to thank the Ph.D. students and postdocs at Eurandom for the many games of fussball and table tennis, watching football and playing poker, and even for our unfortunate attempts at pool. My friends have proven to be a welcome distraction from work, and have helped me to stay motivated. Finally I would like to thank my parents and my sister for their unconditional love.

Peter van de Ven White Plains, October 2011

ACKNOWLEDGEMENTS

CONTENTS

Ac	cknowledgements	\mathbf{V}
1	Introduction1.1Background1.2MaxWeight scheduling1.3Carrier-sense multiple-access1.4Overview of the thesis	1 7 9 16
2	Instability of MaxWeight scheduling2.1Model description2.2Necessary and sufficient stability condition2.3Instability of MaxWeight scheduling2.4Numerical experiments2.5Concluding remarks2.AAuxiliary results	 21 22 23 27 29 30 33
3	Spatial inefficiency of MaxWeight scheduling3.1Model description3.2Instability of MaxWeight Scheduling3.3Stability of region-based scheduling3.4Concluding remarks3.ARemaining proofs	37 38 38 41 46 48
4	Stability and insensitivity4.1Insensitivity of the saturated model4.2A necessary stability condition4.3Stability for full conflict graphs4.4Stability for partial conflict graphs4.5Concluding remarks4.4Auxiliary results and remaining proofs	51 52 54 57 60 62 64
5	Fairness in linear networks5.1Model description5.2Unfairness5.3Achieving fairness5.4Network-average throughput5.5Fair back-off rates in unsaturated networks	 69 70 71 73 77 81

CONTENTS

	5.6 Concluding remarks5.A Proof of Proposition 5.1	82 83
6	Achieving target throughputs	85
	6.1 Model description	86
	6.2 Global invertibility	87
	6.3 Inversion methods	88
	6.4 Special conflict graphs	92
	6.5 Concluding remarks	95
	6.A Remaining proofs	95
7	Optimal tradeoff between exposed and hidden nodes	103
	7.1 Model description and preliminary results	104
	7.2 Main results	107
	7.3 Normalization constant roots	111
	7.4 Optimal sensing range for general networks	113
	7.5 Concluding remarks	117
	7.A Remaining proofs	119
8	Time-slotted CSMA	127
	8.1 Model description	128
	8.2 The saturated regime	129
	8.3 Stability and end-to-end throughput	132
	8.4 Comparing slotted and continuous-time CSMA	133
	8.5 Concluding remarks	135
	8.A Remaining proofs	136
Bibliography		139
Summary		147
Ał	oout the author	149

VIII

Next-generation wireless networks will likely evolve from cellular and small-scale home networks to large, inter-connected networks that form the backbone for lowcost internet access. Such large-scale networks are difficult to evaluate due to the complex spatial and temporal interactions among their users; small networks with few users, in contrast, are relatively well-understood. The study of many-user networks requires models that capture distinct aspects of wireless networks such as interference and the role of medium access control, as well as traffic characteristics and congestion effects. Traditional queueing models are unable to capture the interaction between users, while current models specifically geared towards wireless networks are often limited in scope and network topology, and do not take traffic behavior into account.

In this thesis we develop and examine various mathematical models that capture how users share the wireless medium. We aim to gain a better understanding of wireless networks, and devise schemes to improve their performance. In this chapter we provide a brief introduction to wireless networks, present an overview of the most relevant literature, and summarize the results obtained in this thesis.

1.1 Background

A wireless network can be modeled as a collection of *nodes* (representing users) that can transmit and receive data. Two nodes can be grouped into a transmitter-receiver pair to form a *link*, as shown in Figure 1.1. Here the nodes are represented by circles, while an arrow indicates a link from transmitter to receiver. A node may receive data from different sources, and can transmit towards various destinations. Thus, a node can be associated with multiple links.

A link indicates *potential* data transmission from the transmitting node to the receiver, through the wireless medium. Links can be either active or inactive, depending on whether data is currently being transmitted on that link or not. Let *n* denote the number of links, then the network state can be represented by a vector



Figure 1.1: A wireless network consisting of various nodes and links.

 $\boldsymbol{\omega} = (\omega_1, \omega_2, \dots, \omega_n)$, where ω_i describes the state of link *i* as

$$\omega_i = \begin{cases} 1, & \text{if link } i \text{ is active,} \\ 0, & \text{otherwise.} \end{cases}$$

1.1.1 Interference constraints

Wireless communications are commonly characterized by their broadcast nature, as wireless signals typically propagate in all directions rather than towards the intended receiver of the signal only. As a result, nodes may hear many ongoing transmissions, including those intended for others. In fact, a transmission may not be received correctly if the intended receiver overhears too much conflicting activity. We say in this case that the transmission has suffered a collision due to the *interference* caused by other ongoing transmissions.

Wireless signals are transmitted at a certain power, and the success of a transmission depends on its signal strength as seen by the receiver compared to the strength of the competing transmissions. The strength of a signal decreases with distance, so a wireless network can support multiple simultaneous successful transmissions, but only if the active links are sufficiently far apart.

We assume that all activity conducted by competing links contributes to the interference, and that all interference is treated as noise. In principle this need not be the case since clever coding schemes may mitigate or even completely cancel the adverse 1.1 BACKGROUND

effect of simultaneous transmissions on each other, cf. [48]. However, such coding schemes are difficult to implement, and are not very common in practice.

Whether or not a transmission is successful depends on many different factors such as fading, shadowing and capture effects, which are difficult to determine exactly. In the research literature, various models are presented that describe in detail when a transmission is successful. We will focus on the *physical model* and the *protocol model* [30].

• **Physical model.** We denote by P_i the power at which the signal over link *i* is transmitted, and G_{ij} represents the fraction of signal strength remaining (path loss) after traveling from the transmitter of link *i* to the receiver of link *j*. Thus, the receiver of link *i* overhears a signal of strength $P_jG_{ji}\omega_j$ coming from the transmitter of link *j*. In the physical model the success of a transmission is determined by the ratio of the strength of the transmission signal at the receiver and the background noise *N* plus noise it receives from other transmissions. A transmission on link *i* is considered successful if and only if the Signal to Interference-plus-Noise Ratio (SINR) is above a certain threshold ξ :

$$SINR_{i} = \frac{P_{i}G_{ii}}{N + \sum_{j \neq i} P_{j}G_{ji}\omega_{j}} \ge \xi.$$
(1.1)

3

Denote by X_i and Y_i the locations of the transmitter and receiver of link *i*, respectively. A common assumption is that signal strength attenuates according to a power law, i.e., $G_{ij} = ||X_i - Y_j||^{-\gamma}$, with $|| \cdot ||$ the Euclidian distance and γ the *path loss exponent*. This exponent depends on the environment, and is usually assumed to take values between $\gamma = 2$ (free space) and $\gamma = 4$ (lossy environments).

• **Protocol model.** According to this model, a transmission on link *i* is successful if and only if

$$||X_j - Y_i|| \ge (1 + \Delta)||X_i - Y_i||, \quad \forall j \neq i: \ \omega_j = 1,$$
(1.2)

for some *guard zone* $\Delta > 0$. Essentially, (1.2) says that all links within a certain distance of the receiver have to be inactive in order for a transmission to be successful; the required distance is determined by the guard zone. If all links have the same length *d* (distance between transmitter and receiver), then (1.2) gives rise to an interference range $\eta = (1 + \Delta)d$ centered around the receiving node of a link. A transmission over this link will be successful if and only if no nodes within the interference range are transmitting.

If, depending on the choice of model, (1.1) or (1.2) is satisfied for every active link *i*, all ongoing transmissions are successful. We say that such a state $\omega \in \{0, 1\}^n$ is collision-free, and denote by $\Psi \subseteq \{0, 1\}^n$ the set of all collision-free states.

The physical model gives a more detailed description of the wireless network compared to the protocol model, as it factors in transmission power and signal attenuation, rather than just the distance between nodes. In regimes in which only one or a few links significantly contribute to the interference, the physical model and protocol model are very similar. This is the case for instance if nodes are far apart (sparse networks) or if the signal strength (in the physical model) decreases rapidly with distance (γ large).

To decide whether a state $\omega \in \Psi$ remains collision-free after activating some link *i*, we have to compute the mutual interference between link *i* and the links already active. For the physical model, a transmission on link *i* will affect the entire network, whereas for the protocol model only links *j* such that $||X_i - Y_j|| \le (1 + \Delta)||X_j - Y_j||$ and $||X_j - Y_i|| \le (1 + \Delta)||X_i - Y_i||$ have to be inspected for activity. We say that feasibility for the protocol model can be verified locally, as opposed to globally for the physical model.

1.1.2 Capacity region

An important measure of the quality of a link is the *throughput* θ_i , defined as the expected long-term number of successful packet transmissions over link *i* per time unit. We denote by $\theta = (\theta_1, \theta_2, ..., \theta_n)$ the throughput vector that describes the throughput of all links, and we are interested in the capacity region *C* of the network, defined as all possible values that the throughput vector can take, given the network structure.

The throughput vector is restricted by the interference constraints ((1.1) or (1.2), for instance), and can be attained if and only if there exists some time-sharing of the collision-free states that yields these throughputs. Assuming that transmissions are completed at unit rate, the capacity region of the network can be written as the convex hull of Ψ :

$$C = \operatorname{conv}(\Psi) = \{ \boldsymbol{\theta} \in [0,1]^n \mid \boldsymbol{\theta} = \sum_{\boldsymbol{\omega} \in \Psi} \alpha(\boldsymbol{\omega})\boldsymbol{\omega}, \sum_{\boldsymbol{\omega} \in \Psi} \alpha(\boldsymbol{\omega}) = 1, \ \alpha(\boldsymbol{\omega}) \ge 0 \ \forall \boldsymbol{\omega} \in \Psi \}.$$
(1.3)

The rate at which packets are transmitted may vary between links, depending on packet length, transmission power, and channel state among other things. Denote by R_i the transmission rate across link *i*, defining the expected number of packets that are transmitted per time unit if link *i* is active. The capacity region in this case is similar to (1.3), only with the activity of all links weighted with their respective transmission rates. Transmission rates may also fluctuate over time, due to changes in the channel conditions. Assuming that the feasible transmission rates evolve in a Markovian fashion over a finite number of channel states, the capacity region is given by a weighted average over the capacity region associated with each channel state (see [67]).

The above description is limited to a single-hop capacity region, where all traffic is transmitted directly from source to destination. Alternatively one may look at a multihop capacity region, by allowing intermediate nodes to forward messages intended for others. The advantage of multi-hop communication is that by routing traffic through a series of nearby nodes, the transmit power required for each individual transmission is reduced, which may increase spatial reuse. Moreover, the use of intermediate nodes allows for communication over larger distances than would be possible otherwise. The multi-hop capacity of a network can then be computed by varying the transmit power and packet routing. This approach is taken in [30], where it is investigated how the multi-hop network capacity scales with the number of nodes, under the assumption of equal throughputs for every source-destination pair in the network. It is shown that the throughput for every source-destination pair scales like $m^{-3/2}$ as $m \to \infty$.

1.1 BACKGROUN	D
---------------	---

This work has generated a lot of interest in scaling laws for wireless network capacity under various assumptions; see [103] for an overview.

1.1.3 Traffic modeling

There exists a variety of approaches for modeling the arrival of traffic into the network. The bulk of this thesis is concerned with the *saturated* model, where links are assumed to always have packets available for transmission. This model represents congested traffic conditions as well as networks that operate under certain networklayer protocols that ensure that links are never starved.

The second traffic scenario under consideration is the *unsaturated* model, where packets arrive at the links according to some external arrival process, see Figure 1.2. Packets are temporarily stored in a buffer at the corresponding link pending transmission. In the unsaturated scenario buffers may occasionally be empty, during which time the corresponding link cannot activate. The number of packets stored in a buffer is called the backlog or queue length of a link. We assume that packets arrive according to a renewal process, with λ_i the packet arrival rate (reciprocal of the expected inter-arrival time) at link *i*, and we write $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_n)$. In Chapter 4 we analyze this model under the assumption that packets leave the system immediately once transmitted (Figure 1.2(a)), and in Chapter 8 we consider a multi-hop scenario where packets may be routed between nodes (Figure 1.2(b)).



Figure 1.2: Single-hop and multi-hop unsaturated networks.

In Chapters 2 and 3 we consider a traffic model where the collection of links evolves over time, so-called *flow-level dynamics*. New transmitter-receiver pairs form flows that arrive into the system at random times and locations with some finite number of packets to be transmitted. A flow will leave the system once it has transmitted all its packets. This is illustrated in Figure 1.3, which shows three snapshots of the network evolution. Alternatively, one may consider a hybrid traffic model that combines persistent flows and short-lived flows [54, 55].

5



Figure 1.3: An illustration of flow-level dynamics.

The notion of a capacity region as discussed in Section 1.1.2 is predicated on the assumption of a fixed set of links, and does not readily apply in the case with flow-level dynamics. However, one could define the throughput to be the aggregate transmission rate over all users, so as to arrive at a scalar quantity that measures the total number of packets transmitted.

1.1.4 Medium access control

We have seen that transmissions are subject to certain interference constraints: Only certain subsets of links can be activated simultaneously without giving rise to collisions. Since collisions degrade the network performance, it is essential to devise algorithms that regulate the link activity to reduce interference. Many such *medium access control algorithms* exist, with different implementations and varying degrees of efficacy in preventing collisions. We consider both discrete-time algorithms, where link activity can be changed at the beginning of each time slot t = 0, 1, ..., and continuous-time algorithms where the set of active links can be modified at any time instant $t \ge 0$. Throughout this thesis it will be clear from the context whether t is discrete or continuous.

We distinguish between two classes of access schemes: *scheduled-access algorithms* (discrete-time only) and *random-access algorithms* (both continuous-time and discrete-time). Random-access algorithms form a class of distributed, randomized access schemes, where links decide for themselves when to activate, based on local information only. Due to their localized nature, and since link activity is based to some extent on chance, random-access algorithms may not entirely preclude collisions. It is possible to synchronize all links using message passing algorithms, although this is not required. Consequently, for many random-access algorithms both slotted and non-slotted versions exist, such as the Aloha algorithm [2, 73] and the Carrier-Sense Multiple-Access (CSMA) algorithm [19, 44].

Scheduled-access algorithms implement a time-slotted mechanism, where in each slot a new set of links is selected for transmission. Because of the additional coordination among links, scheduled-access algorithms typically satisfy the interference constraints. Scheduled-access algorithms can be implemented both in a centralized and a distributed way. The former employs a centralized entity that controls the behavior of all links, while in a distributed implementation links decide for themselves

1.2 MAXWEIGHT SCHEDULING

when to activate, based on local information and message passing.

In this thesis we focus on two medium access control algorithms. In Section 1.2 we describe the MaxWeight scheduling algorithm, a centralized mechanism that schedules transmissions so as to maximize a certain weight. Section 1.3 discusses the random-access CSMA algorithm, under which links activate and deactivate autonomously and asynchronously.

1.1.5 Stability region

In Section 1.1.4 we have seen that there exists a wide range of algorithms for sharing access to the wireless medium. These algorithms vary in implementation complexity and performance. In the case of saturated traffic conditions we see that different algorithms may result in markedly different throughput vectors.

Throughput is also an important performance measure in the unsaturated case, but additionally we can ask ourselves whether the network is *stable* under a particular algorithm and given certain traffic conditions. Stability of the network roughly means that the throughput of each link is equal to its arrival rate, so it is not overloaded. In contrast, the throughput of an unstable link is lower than the arrival rate. We consider two definitions of stability: (i) the queues at the various links empty infinitely often with finite expected time (positive recurrence in case the queue length process is a Markov process); and (ii) rate stability, i.e., the departure rate equals the packet arrival rate. Note that definition (i) is stronger than (ii), because a rate stable system does not necessarily empty in finite expected time.

The *stability region* of a scheduler is defined as the set of all arrival rate vectors that yield a stable network. The stability region of a specific policy should be distinguished from the capacity region of the entire network. Naturally, the stability region of a scheduler is always contained in the capacity region of the network, since the latter marks the physical limits of the network transport capacity. When the stability region of a scheduler is identical to (the interior of) the capacity region, we say that this scheduler is throughput-optimal or maximum stable. Ideally we would like to find throughput-optimal schedulers that are applicable in a wide variety of scenarios, without prior knowledge on the network parameters.

1.2 MaxWeight scheduling

The MaxWeight scheduling algorithm is a time-slotted algorithm that has gained immense popularity as a powerful concept for achieving maximum throughput and queue stability in a wide variety of scenarios. It works in a time-slotted fashion, and schedules a collision-free subset of flows (links) for transmission in each slot. Denote by $R_i(t)$ the number of packets that flow *i* could transmit if selected for transmission in time slot *t*. Let $Q_i(t)$ denote the queue length of flow *i* at the beginning of slot *t*, then the MaxWeight scheduling algorithm selects a set of flows so as to maximize the aggregate product of queue length and feasible transmission rate:

$$\underset{\omega \in \Psi}{\operatorname{arg\,max}} \sum_{i=1}^{n} Q_i(t) R_i(t) \omega_i.$$
(1.4)

7

In a pivotal paper [88], a MaxWeight scheduling policy is considered for throughput maximization in multi-hop wireless networks, where only those subsets of links may be activated simultaneously that satisfy the interference constraints, see also [37]. In [89] a MaxWeight policy for allocating a server among several parallel queues with time-varying connectivity is described.

Broadening the latter framework, MaxWeight-type policies were developed for power control and scheduling of wireless channels with rate variations, see for instance [4, 24, 66, 67]. Extending the scope further, in [23, 68, 82, 83] algorithms were devised for joint congestion control, routing and scheduling based on MaxWeight principles. The powerful properties of MaxWeight-type policies have emerged as one of the central paradigms in the broader realm of cross-layer control and resource allocation in wireless networks, see [28] for a comprehensive overview.

MaxWeight-type algorithms have also been proposed for throughput maximization in input-queued switches, where only certain subsets of input-output pairs (e.g., matchings) may be simultaneously connected because of compatibility constraints, see for instance [58, 59]. Extensive background material on MaxWeight policies is contained in [61]. Crucial heavy-traffic results for MaxWeight algorithms were obtained in [81].

A particularly appealing feature is that MaxWeight policies only need information on the current backlogs and instantaneous service rates, and do not rely on any explicit knowledge of the rate distributions or the traffic parameters. On the downside, finding the maximum weight subset is often a challenging problem and potentially NP-hard. This is exacerbated in a distributed setting, where message passing and exchange of backlog information create a substantial communication overhead in addition to the computational burden. This issue is especially pertinent as the maximum weight problem generally needs to be solved at a very high pace, commensurate with the fast time scale on which scheduling algorithms tend to operate. In order to address this issue, it was shown in [15, 24, 87] that randomized policies involve less stringent requirements and yet suffice for achieving maximum stability. In addition, several authors have considered algorithms that solve the maximum weight problem in some approximate sense, and quantified the resulting penalty in guaranteed throughput, see for instance [51, 77, 78, 100, 101].

1.2.1 Flow-level dynamics

As mentioned above, MaxWeight-type policies have been shown to achieve maximum stability under fairly mild assumptions. A fundamental premise however is that the network consists of a fixed set of queues with stationary ergodic traffic processes. In reality, the number of users in the wireless network dynamically varies, as sessions eventually end, while new sessions occasionally start. In many situations the assumption of a fixed set of queues is still a reasonable modeling convention since the scheduling actions and packet-level queue dynamics tend to occur on a very fast time scale, on which the population of active sessions evolves only slowly. In other cases, however, sessions may be relatively short-lived, and the above time-scale separation argument does not apply. The impact of flow-level dynamics over longer time scales is particularly relevant in assessing stability properties, as the notion of stability only has strict meaning over infinite time horizons.

Motivated by the above observations, we examine the stability properties of Max-

1.3 CARRIER-SENSE MULTIPLE-ACCESS

Weight scheduling policies in a scenario with flow-level dynamics. We demonstrate in Chapters 2 and 3 that the maximum stability guarantees are no longer valid in this case. For transparency, we focus on a point-to-point shared wireless downlink channel with rate variations in Chapter 2, and do not consider multi-hop scenarios. In Chapter 3 we show that rate variations are not necessary for the instability to arise, and we show that MaxWeight scheduling is not throughput-optimal in a spatial setting with fixed transmission rates either.

9

The intuitive explanation of the instability encountered in Chapter 2 is that Max-Weight policies tend to favor flows with large backlogs, even when their service rates are not particularly favorable, and thus the rate variations of flows with smaller backlogs are not fully exploited. In Chapter 3 we see that MaxWeight policies may constantly get diverted to arriving flows, while neglecting the opportunity to exploit higher spatial reuse patterns involving a persistently growing number of flows with relatively small remaining backlogs, so the opposing effect is never triggered.

Note that flows with large backlogs are also favored in the absence of any flow-level dynamics. In that case, however, the phenomenon cannot persist since the flows with smaller backlogs will build larger queues and gradually start receiving more service, creating a counteracting force.

It is worth drawing a distinction with [52, 63] that show the stability of joint scheduling and congestion control algorithms in the presence of flow-level dynamics without relying on the conventional simplifying time scale separation argument. The main difference with Chapters 2 and 3 lies in the fact that in these studies the set of flow routes is fixed and that scheduling operates at a class level. Inspection of the results in Chapters 2 and 3 suggests that conventional forms of congestion control would not prevent the kind of instability phenomenon that we observe. In other words, the root cause for the instability appears not to be the lack of congestion control, but the fact that the rate variations are not maximally exploited in the presence of flow-level dynamics.

In the spatial setting of Chapter 3, the possibly unbounded number of flow locations greatly exacerbates the computational complexity of solving the maximumweight problem noted earlier. However, in the analysis we assume that the maximumweight problem itself is solved to optimality in each time slot. Thus the instability of MaxWeight policies as discussed above is entirely disjoint from the throughput penalty which may result from solving the maximum-weight problem only approximately as considered for example in [51, 77, 78, 101].

1.3 Carrier-sense multiple-access

Random-access algorithms form a distributed alternative for centralized mechanisms such as MaxWeight scheduling. Nodes using a random-access algorithm decide for themselves when to transmit, based only on local information. The first such random-access algorithm was Aloha [2]. After finishing a transmission, nodes using this algorithm will remain silent for some random time, before activating again. This so-called *back-off mechanism* reduces simultaneous activity of nearby links, and hence helps to prevent, although not preclude, collisions. The back-off mechanism is implemented by drawing some random back-off time, and then counting down at unit rate; a new

transmission is started when the back-off timer expires. The CSMA algorithm refines Aloha, by introducing a so-called *carrier-sensing mechanism* that tells nodes to monitor nearby activity [44]. Nodes in back-off continuously sense their surroundings, and freeze the back-off timer when they observe too much nearby activity. Only when the measured activity drops below a certain threshold, the back-off process continues to count down. This mechanism reduces collisions since it prevents nearby nodes from activating simultaneously. The CSMA algorithm is for instance implemented in the well-known IEEE 802.11 standard [1].

The CSMA algorithm is studied in Chapters 4-8, where we mostly limit ourselves to the case that nodes have at most one destination, i.e., each node is the transmitter of at most one link. Thus we can uniquely associate every link with its transmitting node, and we can modify the notation and terminology introduced earlier in this chapter accordingly. So in the following when we discuss for example the activity (ω_i), throughput (θ_i), position (X_i) and transmit power (P_i) of node *i*, we refer to the corresponding variables of the link to which node *i* is the transmitter. The transition from links to nodes is done to simplify notation and terminology only, and all our results hold for the more general model where nodes may have multiple destinations. In fact, in Chapter 7 we consider a model where nodes are associated with two receivers.

Similar to the discussion on interference constraints in Section 1.1.1, we may employ various models to decide whether the carrier-sensing mechanism of a node is triggered given a certain configuration of active links. According to the physical model, the carrier-sensing mechanism of link i is triggered if

$$N + \sum_{j \neq i} P_j ||X_i - X_j||^{-\gamma} \omega_j \ge \zeta,$$
(1.5)

that is, if the aggregate noise and interference level exceeds some *carrier-sensing threshold* ζ . The protocol model gives rise to a certain *carrier-sensing range* β such that the carrier-sensing mechanism is triggered if at least one node within distance β is transmitting, i.e.,

$$||X_i - X_j|| \le \beta, \quad \text{for some } j \ne i : \omega_j = 1. \tag{1.6}$$

This translates into an undirected *conflict graph*, where the vertices of the graph represent the links of the network, and two links share an edge if and only if their transmitters are within sensing range from each other, see Figure 1.4.

1.3.1 Feasible states and collisions

The carrier-sensing mechanism restricts the possible activity states that the network can take since (1.5) or (1.6) has to be satisfied in order for a node *i* to activate. We denote by $\Omega \subseteq \{0,1\}^n$ the set of *feasible states* according to (1.5) or (1.6), i.e., all states that can be reached under CSMA.For the protocol model, the set of feasible states corresponds to the incidence vectors of all independent sets of the conflict graph. Recently it was shown that one can implement an interference range even for the physical model [26]. This is done by modifying the carrier-sensing mechanism to monitor changes in the received power rather than the instantaneous power only, and using these differentials to compute the distance to all active nodes.

1.3 CARRIER-SENSE MULTIPLE-ACCESS



Figure 1.4: Constructing a conflict graph.

In general, Ψ and Ω are different, and neither set necessarily contains the other. Thus there may be collision-free states that are not feasible, as well as feasible states that are not collision-free. These two types of states are related to the concept of *exposed nodes* and *hidden nodes*, which are discussed in detail in Chapter 7.

Feasible states that are not collision-free correspond to states where one or more collisions occur. The CSMA algorithm does not necessarily completely preclude collisions, since the carrier-sensing is done by the transmitting node, while collisions occur at the receiving end. Due to this information asymmetry, the transmitter is not aware of the exact interference that the receiving node is subjected to. However, if the carrier-sensing mechanism is configured in a sufficiently conservative manner, we can completely exclude the possibility of collisions, i.e., we have that $\Omega \subseteq \Psi$. This is done by by choosing a small sensing threshold ζ (physical model) or setting the sensing range β sufficiently large (protocol model). Note that by doing so we may also eliminate some collision-free states, effectively reducing the network capacity.

In recent years this tradeoff between network capacity and collision reduction has received much attention [50, 57, 104, 107]. Most of these analytic studies assume that the activity of nodes and their back-off processes are independent, which greatly simplifies the analysis. However, the interaction between nodes has a large impact on the performance of the network. The tradeoff between preventing collisions and spatial reuse is the subject of Chapter 7, where we do take this interaction into account, by keeping track of the activity of nodes over time.

1.3.2 CSMA model

We consider a network of *n* nodes sharing a wireless medium according to a CSMAtype protocol. The network is described by an undirected conflict graph (*V*, *E*), where the set of vertices $V = \{1, ..., n\}$ represents the nodes of the network and the set

11

of edges $E \subseteq V \times V$ indicates which pairs of nodes cannot activate simultaneously. In other words, nodes that are neighbors in the conflict graph are prevented from simultaneous activity by the carrier-sensing mechanism. An inactive node is said to be blocked whenever any of its neighbors is active, and unblocked otherwise. We assume for now that the carrier-sensing mechanism is configured such that collisions are completely precluded; this assumption is relaxed in Chapter 7.

Consider a scenario where nodes are saturated, i.e., always have packets to transmit. The transmission times of node *i* are independent and exponentially distributed with mean $1/\mu_i$. When node *i* is blocked, it remains silent until all its neighbors are inactive, at which point it tries to activate after an exponentially distributed (back-off) time with mean $1/\nu_i$. Node *i* activates if it is still unblocked when the back-off timer runs out. If a node finds itself blocked when the back-off timer expires, it waits until all neighboring nodes become inactive once more and then repeats the back-off procedure. Equivalently, we could think of the potential activation epochs of a node as occurring according to a Poisson process, and actual transmission periods starting whenever a potential activation event occurs while the node is unblocked. For conciseness, denote $\sigma_i = \nu_i/\mu_i$.

The set Ω of all feasible joint activity states of the network in this case corresponds to the incidence vectors of all independent sets of the conflict graph. By the assumption that all collisions are precluded, we have $\Omega \subseteq \Psi$. Let the network state at time t be denoted by $\mathbf{Y}(t) = (Y_1(t), Y_2(t), \dots, Y_n(t)) \in \Omega$, with $Y_i(t)$ indicating whether node iis active at time t ($Y_i(t) = 1$) or not ($Y_i(t) = 0$). Then $\{\mathbf{Y}(t)\}_{t\geq 0}$ is a Markov process which is fully specified by the state space Ω and the transition rates

$$r(\omega, \omega') = \begin{cases} \nu_i, & \text{if } \omega' = \omega + \mathbf{e}_i \in \Omega, \\ \mu_i, & \text{if } \omega' = \omega - \mathbf{e}_i \in \Omega, \\ 0, & \text{otherwise.} \end{cases}$$
(1.7)

Here \mathbf{e}_i denotes the vector of length *n* with all zeros except for a 1 at position *i*.

Since $\mathbf{Y}(t)$ is reversible (see [11]), the following product-form stationary distribution π exists:

$$\pi(\boldsymbol{\omega}) = \lim_{t \to \infty} \mathbb{P}(\mathbf{Y}(t) = \boldsymbol{\omega}) = \begin{cases} Z^{-1} \prod_{i=1}^{n} \sigma_{i}^{\omega_{i}}, & \text{if } \boldsymbol{\omega} \in \Omega, \\ 0, & \text{otherwise,} \end{cases}$$
(1.8)

where

$$Z = \sum_{\omega \in \Omega} \prod_{i=1}^{n} \sigma_i^{\omega_i}$$
(1.9)

is the normalization constant that makes π a probability measure. This result is well known in the context of wireless networks, see e.g. [11, 17, 20, 98]. Chapter 4 describes how this result can be extended to general back-off times and transmission durations.

We are interested in the long-term behavior of the network, characterized by the throughput vector θ . As active nodes finish their transmissions at rate μ_i , and all transmissions are successful, we have that

$$\theta_i = \mu_i \sum_{\omega \in \Omega} \pi(\omega) \omega_i. \tag{1.10}$$

1.3 CARRIER-SENSE MULTIPLE-ACCESS

This closed-form expression for the throughput allows for a detailed analysis of the network behavior; this was first done in the 1980s in the context of packet-radio networks [11, 12, 42, 71]. CSMA-type models with arbitrary conflict graphs were first pursued in the context of IEEE 802.11 systems in [98], and further studied in that setting in [19, 21, 22], with several extensions and refinements in [20, 27, 60, 79].

13

In [98] three nodes on a line that only block their direct neighbors are considered. It is shown that the middle node is starved when the back-off rates of all three nodes increase. Such unfairness has been studied for more general networks in [17, 20, 22], and is the subject of Chapter 5 of the present thesis.

Although the representation of the IEEE 802.11 back-off mechanism in the CSMA model is far less detailed than in the landmark work [7], the general conflict graph offers greater versatility and covers a broad range of topologies. Experimental results in [49] demonstrate that these models, while idealized, provide throughput estimates that match remarkably well with measurements in actual IEEE 802.11 systems.

1.3.3 Unsaturated CSMA model

The CSMA model described in Section 1.3.2 focuses on a saturated scenario where nodes always have packets pending for transmission. Alternatively we may consider a network using CSMA in unsaturated traffic conditions, giving rise to queueing dynamics. In particular, the buffers may empty from time to time, and nodes will refrain from competition for the medium during these periods. The resulting interaction between the activity states and the buffer contents of the various nodes gives rise to quite intricate behavior. In particular, the queueing dynamics entail high-dimensional stochastic processes with infinite state spaces, which generally do not admit closed-form expressions for the stationary distribution. Even just establishing the existence of a stationary distribution, i.e., obtaining the stability conditions, is generally a challenging problem, and may often be about as hard as determining the entire joint distribution of the buffer contents.

Unsaturated CSMA models have received little attention in the research literature due to their complexity. In [17, 31] a linear multi-hop wireless network is considered. The end-to-end throughput of a three-node network is computed in [17], and [31] focuses on how to improve the performance of the network by altering the back-off process.

In this thesis we discuss unsaturated CSMA in Chapters 4 and 8. Since a closedform expression for the throughput similar to (1.10) is not available for unsaturated CSMA networks, we instead aim for stability and throughput bounds. Chapter 4 is concerned with the stability region of single-hop CSMA models, in particular in the case of the full conflict graph. In Chapter 8 we study stability and end-to-end throughput of a linear multi-hop network.

1.3.4 Related models

The CMSA model can be interpreted as a special instance of a loss network [38, 40, 41, 84, 105]. Such loss networks were first introduced to study telephone networks, and can be seen as an extension of the classical Erlang loss system [14].

A loss network consists of *J* links (not be be confused with links in the wireless network), where link *j* has C_j circuits j = 1, 2, ..., J. There is a set $\mathcal{R} = \{1, 2, ..., n\}$ of routes, and calls on route $r \in \mathcal{R}$ each use $A_{jr} \in \mathbb{Z}_+$ circuits from link *j*, with \mathbb{Z}_+ the non-negative integers. Calls of type $r \in \mathcal{R}$ arrive according to a Poisson process with rate v_r and have exponentially distributed holding times with mean $1/\mu_r$. If upon arrival of a type-*r* call fewer than A_{jr} circuits are available for any link j = 1, 2, ..., J, the call is rejected.

Denote by $N_r(t)$ the number of calls in progress on route r at time t, and define $\mathbf{N}(t) = (N_r(t), r \in \mathcal{R})$ and $\mathbf{C} = (C_1, C_2, \dots, C_J)$. It is well known (see, e.g., [41]) that the Markov process $\{\mathbf{N}(t)\}_{t\geq 0}$ has a unique stationary distribution

$$\pi(\mathbf{n}) = \lim_{t \to \infty} \mathbb{P}(\mathbf{N}(t) = \mathbf{n}) = Z^{-1} \prod_{r \in \mathcal{R}} \frac{\sigma_r^{n_r}}{n_r!}, \quad \mathbf{n} \in \Omega$$

where

$$\Omega = \{ \mathbf{n} \in \mathbb{Z}_+^{\mathcal{R}} : \mathbf{A}\mathbf{n} \le \mathbf{C} \}$$

with component-wise inequality and $Z = \sum_{\mathbf{n} \in \Omega} \prod_{r \in \Omega} \frac{\sigma_r^{n_r}}{n_r!}$ the normalization constant.

It is readily seen that the CSMA model is in fact a special instance of a loss network, where the call types correspond to the nodes, and the arrival rate v_r is equivalent to the back-off rate. The mean call holding times $1/\mu_r$ are equivalent to the mean packet transmission times. Any CSMA model can be represented as a loss network in multiple ways. For example, consider a CSMA model on some conflict graph (V, E), let $J = |\mathcal{R}| = n$ and choose $C_j = \Delta$, the maximum node degree of the conflict graph. If we then choose

$$A_{jr} = \begin{cases} \Delta, & \text{if } j = r, \\ 1, & \text{if } \{j, r\} \in E, \\ 0, & \text{otherwise,} \end{cases}$$

we see that the resulting loss network is equivalent to the CSMA model. Alternatively, let $C_j = 1$, $|\mathcal{R}| = n$ and J = |E|. Then for

$$A_{jr} = \begin{cases} 1, & \text{if } j \text{ is an edge to } r, \\ 0, & \text{otherwise,} \end{cases}$$

the resulting loss network is again equivalent to the CSMA model.

Despite the extensive literature on loss networks, the application to CSMA models poses new and challenging questions. Traditionally the main focus in loss networks has been on the loss probability, i.e., the probability that a call arriving into the system cannot be accepted due to insufficient capacity at one or more of its required links. This *loss probability* may be written as

$$L_{r} = \sum_{\substack{\boldsymbol{\omega} \in \Omega\\ \boldsymbol{\omega} + \mathbf{e}_{r} \notin \Omega}} \pi(\boldsymbol{\omega}).$$
(1.11)

Evaluating (1.11) is computationally expensive since it requires summing over all possible system states. Thus much effort has gone into designing approximations and establishing asymptotics for the loss probability. The inverse question of choosing the link capacities to attain sufficiently low loss probabilities has also received considerable attention.

1.3 CARRIER-SENSE MULTIPLE-ACCESS

The main performance measure of CSMA models is the throughput (1.10). Although this is related to the loss probability as

$$\theta_r = \nu_r (1 - L_r), \tag{1.12}$$

15

results on loss networks provide little help in the study of CSMA models. For instance, most approximations for the loss probability are designed for the high-capacity regime, so (1.12) cannot be used to obtain easy approximations for the throughput. Moreover, the design questions are different for both models since in loss networks one typically manipulates the link capacities, which is not possible for the CSMA model.

From the connection with loss networks, it is readily seen that the stationary distribution of the joint activity process of the CSMA model is in fact insensitive to the distribution of the transmission times, i.e., the stationary distribution only depends on the mean transmission time. Although loss networks are not insensitive to the interarrival time distribution, we show in Chapter 4 that CSMA models are insensitive to both back-off times and transmission durations. The reason is that the strict equivalence between the CSMA model and loss networks relies on the back-off periods being exponentially distributed. In order to see that, observe that in loss networks the arrival process is not affected by the occupancy state, whereas in the CSMA model the back-off process of a node is suspended when that node is active, and is possibly frozen by the activity of neighboring nodes. In case the back-off periods are exponentially distributed, back-off freezing does not affect the activity process, so the CSMA model is equivalent to a loss network. For generally distributed back-off periods this distinction does become relevant, and no direct analogy with loss networks applies.

Another interesting connection appears when we look at the Markov chain obtained by embedding the Markov process of the CSMA model on transition instants. This Markov chain in fact is equivalent to the Glauber dynamics of the hard-core model [47] from statistical physics. In Section 1.3.5 we describe how the connection is used to design adaptive CSMA algorithms.

1.3.5 Adaptive CSMA

Traditional CSMA assumes that the mean back-off times and transmission durations remain fixed over time. Recently, several clever adaptive CSMA-type algorithms have appeared which achieve throughput-optimality by adjusting the back-off rates over time. In [32, 34], a class of distributed algorithms is proposed, where nodes adjust their back-off rates based on current backlog, which is defined as the difference between arrived traffic and transmitted packets, while [72] suggests to choose the back-off rate to be a certain increasing function of the backlog. In [34] it is shown that these protocols can achieve any throughput vector in the interior of the capacity region (1.3).

The key idea of the algorithm in [34] is to adapt the back-off rates of the nodes according to the difference between arrival rate and throughput. This difference is exactly the gradient associated with a specific convex optimization problem, the solution of which provides stability, if possible to do so at all. In [33, 36, 53] it is shown that the back-off rates prescribed by this algorithm converge. This approach can be used to optimize a utility function of the throughputs, providing for example max-min fairness or maximization of the aggregate throughput.

The approach of [72] is roughly to choose the back-off rates as $\log(q + 1)$, where q is the current backlog. This choice responds very slowly to queue-length increases, and is known to cause long delays. Recently these requirements were relaxed, see [29]. Here it is shown that it is sufficient for maximum stability if the logarithms of the back-off rates behave as $\log(q + 1)/g(q)$, with g(q) strictly increasing and chosen such that $\log(q + 1)/g(q)$ is strictly concave and increasing. It is shown by means of simulation that this choice for the weights leads to lower average delay. In [9] it is shown that even linear weights provide maximum stability, but this was proven under a time-scale separation assumption. A similar approach was taken for the multi-channel case in [10].

The above adaptive algorithms show remarkable performance in terms of throughput, but are reported to cause very long delays. In [56, 75], the specific structure of the conflict graph that arises in wireless networks is exploited to devise CSMA algorithms where the delay *does* scale well with the network size. In [75] this is done by temporarily freezing some nodes, whereas [56] suggests to occasionally shut down and then restart the entire network.

The above references assume an idealized setting without collisions. In [35, 70, 76] different adaptive CSMA algorithms are described, that are throughput-optimal even in a setting with a certain type of collisions. This is done in [70, 76] by considering a discrete-time protocol where each time slot is divided into a control phase and a data phase. The control phase is used to determine a collision-free schedule for the data phase, and the resulting adaptive algorithm is such that no collisions occur during the data phase. The solution proposed in [35] constitutes a continuous-time version. In [43], an algorithm designed to deal with collisions caused by false-negatives of the carrier-sensing mechanism is presented.

While adaptive CSMA achieves throughput-optimality, the case of fixed back-off rates is nevertheless relevant since in practice the adaptation of back-off parameters involves a wide range of non-trivial implementation issues (finite-range precision, communication overhead, information exchange), and hence it is important to gain insight in the achievable performance of non-adaptive algorithms. This is also demonstrated by [45, 65], that implement a version of the adaptive algorithm from [34]. The experiments there show that while adaptive CSMA performs well in certain scenarios, its effectiveness is strongly reduced by various phenomena encountered in practice, such as capture effects and the presence of hidden nodes. For example, hidden-node collisions cause the nodes to become overly aggressive, which may lead to complete starvation of certain other nodes.

1.4 Overview of the thesis

In this thesis we examine various mathematical models in order to improve our understanding of the role of medium access control algorithms in wireless networks. These models exhibit similar qualitative behavior as real-life wireless networks, and can be used to gain insight into various known performance issues, as well as uncover new problems. We focus on the MaxWeight scheduling and CSMA algorithms, both of which are popular mechanisms for regulating node activity and sharing resources in wireless networks. As described earlier, the goal of such algorithms is to allow for

1.4 OVERVIEW OF THE THESIS

simultaneous activity of many links, while restricting the set of active links to certain collision-free subsets. It turns out that MaxWeight scheduling and CSMA, although markedly different, both suffer from performance issues that have the same underlying cause: The algorithms under consideration may consistently schedule unfavorable states, as is illustrated below.

For example, consider a saturated linear CSMA network of three nodes, with nearest neighbor blocking, so only nodes 1 and 3 can be active simultaneously. Assume that all nodes activate with rate $v_i = \sigma$ so that the mean back-off time equals $1/\sigma$. For this small network, the saturation throughputs can be easily computed using (1.10):

$$\theta_1 = \frac{\sigma(1+\sigma)}{1+3\sigma+\sigma^2}, \quad \theta_2 = \frac{\sigma}{1+3\sigma+\sigma^2}, \quad \theta_3 = \frac{\sigma(1+\sigma)}{1+3\sigma+\sigma^2}.$$
 (1.13)

17

As was reported in [98], the throughput is highly unfair, and nodes 1 and 3 receive much better service than the node in the middle. Node 2 can only activate when both outer nodes are silent. As σ increases, this event occurs less frequently, and from (1.13) it is readily seen that node 2 will be completely starved as $\sigma \rightarrow \infty$. In terms of scheduling feasible subsets of nodes, we see that the CSMA algorithm favors the state (1,0,1) over (0,1,0), leading to unfair throughputs.

A similar phenomenon occurs in MaxWeight scheduling, when applied in a setting with flow-level dynamics. Consider the same interference structure as before, only with the nodes replaced by regions. New flows of deterministic size arrive into one of the three regions, and at most one flow per region can be scheduled at any point in time. So the scheduler can choose to select either a flow each from regions 1 and 3 (schedule (1, 0, 1)), or one flow from region 2 (schedule (0, 1, 0)). We assume a fixed transmission rate $R_i(t) \equiv 1$, so that MaxWeight scheduling selects

$$\omega = \begin{cases} (1,0,1), & \text{if } N_1^*(t) + N_3^*(t) \ge N_2^*(t), \\ (0,1,0), & \text{otherwise,} \end{cases}$$

with $N_i^*(t)$ the size of the largest flow in region *i* at time *t*. If new flows in region 2 have unit size, and new flows in regions 1 and 3 have size greater than one, then the MaxWeight scheduling algorithm selects (1,0,1) whenever a new flow arrives in either region 1 or 3, irrespective of the number of flows already present in region 2. This causes the number of flows in region 2 to explode. This behavior is key to the instability of MaxWeight scheduling discussed in Chapter 3.

1.4.1 Instability of MaxWeight scheduling

As already hinted at in the above example, MaxWeight may run into difficulties when confronted with flow-level dynamics. In Chapters 2 and 3 we demonstrate that in the presence of flow-level dynamics the algorithm may no longer be throughput-optimal, and we identify two causes for the instability: (i) failure to fully exploit rate variations; and (ii) spatial inefficiency.

In Chapter 2 we consider the inability of MaxWeight scheduling to exploit rate variations, which can be demonstrated in a single-downlink scenario with varying transmission rates. We identify a simple necessary and sufficient condition for stability, and show that MaxWeight policies may fail to provide maximum stability. The

intuitive explanation is that these policies tend to favor flows with large backlogs, so that the rate variations of flows with smaller backlogs are not fully utilized.

The second cause for instability is studied in Chapter 3, where we consider a spatial setting in which flows arrive at random in some finite space, and multiple flows may be scheduled simultaneously, subject to certain interference constraints. The MaxWeight scheduler tends to serve flows with large backlogs, even when the resulting spatial reuse is not particularly efficient. We show that MaxWeight policies consistently choose inefficient schedules, which may lead to instability.

1.4.2 Insensitivity of the CSMA model

In Section 1.3.4 we explained the connection between CSMA models and loss networks, and argued that CSMA models can be seen as a special instance of loss networks. Loss networks are well known to be insensitive to the distribution of the call holding times, in the sense that the stationary distribution only depends on the mean of the holding time rather than on the entire holding-time distribution, see [41]. It is easily seen that this implies insensitivity of CSMA models to the distribution of the transmission times. Moreover, despite the fact that the insensitivity for loss networks does not extend to interarrival times, we show in Chapter 4 that CSMA models are in fact insensitive to the back-off times. The reason for this is that in CSMA models the back-off process of an active node is suspended, while the arrival process of a blocked route in a loss network continues while blocked.

1.4.3 Stability of random-access networks

In Chapter 4 we also consider the unsaturated model, where packets arrive at each node *i* according to some renewal process with rate λ_i , and buffers may occasionally empty. We are interested in the stability region of the CSMA algorithm.

First we use the corresponding saturation throughput to give a simple sufficient condition for instability:

$$\lambda_i > \theta_i, \quad i = 1, 2, \dots, n,$$

and we show that the converse condition is not sufficient for stability. We then explicitly identify the stability region for the complete conflict graph, and illustrate the difficulties that arise when trying to describe the stability region for partial conflict graphs.

1.4.4 Throughputs and fairness of CSMA

As has been mentioned a few times already, CSMA networks may exhibit severe unfairness, in the sense that some nodes receive consistently higher throughput than others. In Chapter 5 we study this phenomenon in linear networks, and realize strict fairness by choosing certain node-specific back-off rates. We obtain closed-form expressions for the fair back-off rates and the resulting throughputs.

The more general problem of finding the back-off rates that yield a certain throughput vector is addressed in Chapter 6. Let $\gamma = (\gamma_1, ..., \gamma_n)^T \in \mathbb{R}^n_+$ belong to the range Γ of the mapping $\theta : \mathbb{R}^n_+ \to \Gamma$. In [34] it is shown that Γ is equal to the interior of the capacity region *C* of the network. The throughput vector is a highly non-trivial and 1.4 OVERVIEW OF THE THESIS

non-linear function of the back-off rates, and the problem of finding back-off rates that achieve a certain throughput vector can be formalized as finding $\nu_{\theta} = \nu_{\theta}(\gamma)$ that solves

$$\theta(\nu_{\theta}) = \gamma,$$

and hence, we need to study in detail the mapping θ . We show that the throughput function is globally invertible, meaning that for any $\gamma \in \Gamma$ (fair or otherwise) there exists exactly one ν_{θ} that yields γ . In contrast to fairness on a line, we can no longer determine the inverse explicitly. Instead, we present several numerical procedures for calculating the inverse, based on fixed-point iteration and Newton's method.

1.4.5 Carrier-sensing tradeoff

As explained in Section 1.3, the carrier-sensing mechanism of CSMA may not completely preclude collisions. The reason is that whether or not transmissions are successful depends on the noise perceived by the transmitting node, while the carriersensing mechanism is triggered at the transmitting node. We can reduce interference by increasing the carrier-sensing range β , although this also reduces spatial reuse. In Chapter 7 we study this tradeoff in a linear wireless network, for a given interference range η , a conflict graph that arises from the carrier-sensing range β , and uniform back-off rate σ . We express the throughput as a function of various instances of the normalization constant of a linear CSMA model with *i* nodes, as defined in (1.9), and use this to solve for the throughput-optimal value of β . We show that the value of the optimal sensing range depends on the mean back-off times of the nodes.

1.4.6 Time-slotted CSMA

In Chapter 8 we study a time-slotted CSMA algorithm, where all nodes are synchronized and transmissions last exactly one time slot. We consider a linear network and determine the network-aggregate throughput and per-node throughputs under saturation conditions. These are compared to the results obtained for continuous-time CSMA in Chapter 5. We then provide bounds on the end-to-end throughput for both slotted and continuous-time CSMA.

1.4.7 Literature summary

This thesis is largely based on results that have already appeared in the literature, and we proceed to give an overview of the relevant papers. Chapter 2 is based on [90] and Chapter 3 on [94]. In Chapter 4 we present the results from [92] while Chapter 5 follows [91, 96]. The results presented in Chapter 6 were first derived in [95], and Chapters 7 and 8 are based on [93] and [80], respectively.

19

2

INSTABILITY OF MAXWEIGHT SCHEDULING

In Section 1.2 we discussed the celebrated MaxWeight scheduling algorithm, a versatile centralized medium access control mechanism. The popularity of MaxWeight scheduling is due to its ability to provide maximum stability, which is shown to hold in a wide variety of scenarios, but only in case that the system consists of a fixed set of queues with stationary ergodic traffic processes. In reality, the collection of active queues dynamically varies, as flows eventually depart while new flows occasionally start. In the present chapter and in Chapter 3 we will demonstrate that the maximum-stability guarantees of MaxWeight scheduling are no longer valid under flow-level dynamics. In this chapter we focus on a point-to-point shared wireless downlink channel with rate variations.

This chapter is organized as follows. In Section 2.1 we present a detailed model description and in Section 2.2 we derive a simple necessary and sufficient condition for stability in the presence of flow-level dynamics. Section 2.3 establishes that the MaxWeight policy may fail to provide maximum stability by treating specific model instances where the stability conditions are satisfied, yet MaxWeight scheduling does not keep the system stable. In Section 2.4 simulation results are provided that support the analytical findings and in Section 2.5 we make some concluding remarks.

INSTABILITY OF MAXWEIGHT SCHEDULING

2.1 Model description

We consider a single wireless link shared by *K* classes of flows. The system operates in a time-slotted fashion, and in each time slot at most one of the flows can be scheduled for transmission. Denote by $A_k(t)$ the number of class-*k* flows starting in time slot *t*. We assume that $A_k(1), A_k(2), \ldots$ are i.i.d. copies of some random variable A_k with mean $\alpha_k < \infty$. The arrivals are independent both over time and between classes.

Each of the flows generates some finite random amount of traffic. We distinguish between two scenarios for the traffic influx of the various flows: (i) instantaneous traffic bursts; and (ii) gradual traffic streams. In case (i) each flow generates an instantaneous amount of traffic upon arrival to the system. Denote by B_{ki} the size of the *i*-th class-*k* flow upon arrival (in bits). We assume that B_{k1}, B_{k2}, \ldots are i.i.d. copies of some integer random variable B_k with $\mathbb{E}[B_k] < \infty$. The flow sizes upon arrival are independent both over time and between classes.

In case (ii), each flow starts a random finite activity period upon arrival to the system, during which it produces a gradual stream of traffic. Denote by D_{ki} the duration of the activity period of the *i*-th class-*k* flow (in slots). We assume that D_{k1}, D_{k2}, \ldots are i.i.d. copies of some integer random variable D_k with $\mathbb{E}[D_k] < \infty$. Denote by $F_{ki}(t)$ the amount of traffic in bits generated by the *i*-th class-*k* flow in time slot *t*. For notational convenience, we define $F_{ki}(t)$ for all *t*, but its value is only relevant if the *i*-th class-*k* flow is active. We assume that $F_{ki}(1), F_{ki}(2), \ldots$ are i.i.d. copies of some integer random variable F_k with $\mathbb{E}[F_k] < \infty$, the $F_{ki}(1)$ are independent from the D_{ki} , and that the traffic processes are independent among the various flows. Denote by $B_{ki} = \sum_{t=S_{ki}}^{S_{ki}+D_{ki}-1} F_{ki}(t)$ the total amount of traffic generated by the *i*-th class-*k* flow, with S_{ki} denoting its arrival time. By the above assumptions, B_{k1}, B_{k2}, \ldots are i.i.d. copies of an integer random variable B_k with mean $\mathbb{E}[B_k] = \mathbb{E}[D_k]\mathbb{E}[F_k] < \infty$.

Note that scenario (i) may be interpreted as a special case of scenario (ii) with $D_k \equiv 1$ and $F_k \equiv B_k$. For economy of notation, however, it is useful to classify scenario (i) as a separate case. In both scenarios, traffic may only start to be served in the next slot after it arrives. Flows leave the system as soon as all their bits have been transmitted (and no further bits are due to arrive in the case of gradual traffic streams). During the period between its arrival and departure, a flow is said to be present.

The feasible transmission rates of the various flows vary over time as a result of fading. Denote by $R_{ki}(t)$ the feasible transmission rate (in bits) of the *i*-th class-*k* flow if selected for transmission in time slot *t*. For notational convenience, we define $R_{ki}(t)$ for all *t*, but its value is only relevant if the *i*-th class-*k* flow is actually present in the system. We assume that $R_{ki}(1), R_{ki}(2), \ldots$ are i.i.d. copies of some integer, positive random variable R_k , and that the feasible transmission rates are independent among the various flows. Define $R_k^{max} = \sup\{r : \mathbb{P}(R_k = r) > 0\}$ as the maximum possible value of the transmission rate of class-*k* flows (possibly $R_k^{max} = \infty$).

The flow arrivals, sizes and feasible transmissions rates are extraneous, while we can choose which flow to schedule in each time slot. Let us say that in time slot t a flow of class k(t) with a residual size of l(t) bits is served at rate r(t), with the convention that k(t) = l(t) = r(t) = 0 in case no flow gets scheduled in time slot t at all.

The evolution of the system over time in case of instantaneous traffic can be described by a vector $N(t) = (N_1(t), \dots, N_K(t))$, with $N_k(t) = (N_1^k(t), N_2^k(t), \dots)$ and

2.2 NECESSARY AND SUFFICIENT STABILITY CONDITION

 $N_l^k(t)$ representing the number of class-*k* flows in the system with a residual size of *l* bits at the beginning of slot *t*. Observe that

23

$$N_{l}^{k}(t+1) = N_{l}^{k}(t) + A_{l}^{k}(t) - \mathbb{1}_{\{k(t)=k, l(t)=l\}} + \mathbb{1}_{\{k(t)=k, l(t)=l+r(t)\}},$$

with $A_l^k(t)$ denoting the number of class-*k* flows arriving at time *t* with a size of exactly *l* bits. It is easily verified that the process N(t) is a Markov chain. A similar description of the system evolution for gradual traffic is provided in the proof of Theorem 2.2.

Define $\rho_k = \alpha_k \tau_k$, and $\rho = \sum_{k=1}^{K} \rho_k$, with $\tau_k = \mathbb{E}[[B_k/R_k^{\max}]]$ when $R_k^{\max} < \infty$ and $\tau_k = 1$ when $R_k^{\max} = \infty$. Thus τ_k represents the expected number of slots required for the service of a class-*k* flow when served at rate R_k^{\max} .

2.2 Necessary and sufficient stability condition

In this section we first establish a simple necessary condition for stability to be achievable, and then proceed to show that this is in fact also (nearly) sufficient. The system is said to be stable if the Markov chain that describes the state of all present flows is positive recurrent.

PROPOSITION 2.1. *The condition* $\rho \leq 1$ *is necessary for stability.*

PROOF. The expected number of slots required for the service of an arbitrary class-*k* flow is bounded from below by τ_k . Thus the rate at which class-*k* work enters the system is bounded from below by $\rho_k = \alpha_k \tau_k$, and the total rate at which work arrives is bounded from below by $\rho = \sum_{k=1}^{K} \rho_k$. The latter quantity may not exceed one in order for stability to be achievable.

We proceed to show that the above condition is also (nearly) sufficient for stability to be achievable. This may be intuitively explained as follows. With a dynamic population of flows, there will always be a flow that has the maximum possible feasible rate with high probability when there are sufficiently many flows present in the system. In other words, whenever a flow gets selected for transmission, it can be served at the maximum possible rate with high probability. Thus the expected number of slots required for the service of an arbitrary class-*k* flow can be brought arbitrarily close to τ_k , so that the system can be stabilized for values of ρ arbitrarily close to 1.

Evidently, the above explanation only provides heuristic arguments and does not account for several subtle yet critical issues. However, the intuitive insight offers useful guidance for the construction of a Lyapunov function that serves as the basis of a rigorous proof of the propositions presented below.

We distinguish between the two traffic scenarios described in the previous section. As mentioned earlier, the scenario with instantaneous traffic bursts may be interpreted as a special case of that with gradual traffic streams. For transparency, however, we provide a separate treatment which introduces the key concepts while avoiding some of the additional complexity that arises in the general case.

THEOREM 2.1. For any $\rho < 1$, there exists a scheduling strategy that achieves stability in case of instantaneous traffic.

INSTABILITY OF MAXWEIGHT SCHEDULING

PROOF. We first introduce several constants that will be used. Let $\epsilon = \frac{1}{2}(1-\rho)/(K+1) > 0$ and define $Z_k = \min\{R_k^{\max}, \alpha_k \mathbb{E}[B_k]/\epsilon\}, \eta_k := \mathbb{P}(R_k \ge Z_k) > 0, \eta := \min_{k=1,\dots,K} \eta_K$ and $N_{\epsilon,\eta} := \min\{m : (1-\eta)^m \le \epsilon\}$. Denote $\theta_k = [\mathbb{P}(R_k > 0)]^{-1} < \infty$, let $L_k = \min\{l : \sum_{i=l+1}^{\infty} i\mathbb{P}(B_k = i) \le \epsilon/(\alpha_k \theta_k)\}$, and observe that $L_k < \infty$ since $\mathbb{E}[B_k] < \infty$.

We consider a scheduling strategy with the following property: it serves a class-k flow that either (i) has a feasible transmission rate Z_k or higher or (ii) has a residual size L_k or larger and a positive feasible transmission rate, whenever possible. Ties are broken arbitrarily. In order to prove stability, we have to show that the Markov chain N(t) is positive recurrent.

Define the Lyapunov function:

$$V(n) = \sum_{k=1}^{K} \left(\sum_{i=1}^{L_k} n_i^k \left\lceil \frac{i}{Z_k} \right\rceil + \theta_k \sum_{i=L_k+1}^{\infty} i n_i^k \right),$$

with $n = (n_1, ..., n_K)$ and $n_k = (n_1^k, n_2^k, ...)$.

The function V(n) provides a measure for the total amount of work in the system in terms of the total number of slots required for the service of all currently present flows, assuming that class-*k* flows of residual size no larger than L_k are always served at rate Z_k , while class-*k* flows of residual size of at most L_k are served at rate $\theta_k^{-1} = \mathbb{P}(R_k > 0)$.

We can write the drift as

$$V(N(t+1)) - V(N(t)) = \sum_{k=1}^{K} I_k(t) - D(t),$$

with

$$I_{k}(t) = \sum_{i=1}^{L_{k}} \left[\frac{i}{Z_{k}} \right] A_{i}^{k}(t) + \theta_{k} \sum_{i=L_{k}+1}^{\infty} i A_{i}^{k}(t),$$
(2.1)

reflecting the increase in the workload due to the arrival of class-k flows, and

$$D(t) = \left[\frac{l(t)}{Z_{k(t)}}\right] \mathbb{1}_{\{1 \le l(t) \le L_{k(t)}\}} + \theta_{k(t)} l(t) \mathbb{1}_{\{l(t) > L_{k(t)}\}} - \left[\frac{l(t) - r(t)}{Z_{k(t)}}\right] \mathbb{1}_{\{1 \le l(t) - r(t) \le L_{k(t)}\}} - \theta_{k(t)} (l(t) - r(t)) \mathbb{1}_{\{l(t) - r(t) > L_{k(t)}\}}$$

$$(2.2)$$

representing the decrease in the workload due to the service of flows. The conditional drift may then be written as:

$$\mathbb{E}[V(N(t+1)) - V(N(t)) \mid N(t) = n] = \sum_{k=1}^{K} \mathbb{E}[I_k(t)] - \mathbb{E}[D(t) \mid N(t) = n].$$
(2.3)

Define

$$C = \{n \mid \sum_{k=1}^{K} n_k < N_{\epsilon,\eta} \text{ and } \sum_{k=1}^{K} s_k = 0\}.$$

It may be shown that

 $\mathbb{E}[I_k(t)] \le \rho_k + 2\epsilon, \tag{2.4}$

$$\mathbb{E}[D(t) \mid N(t) = n] \ge 1 - \epsilon, \quad n \notin C.$$
(2.5)

Both (2.4) and (2.5) are derived in Lemma 2.1, which is presented and proven in Appendix 2.A.

25

Combining Equations (2.3)-(2.5) we get

$$\mathbb{E}[V(N(t+1)) - V(N(t)) \mid N(t) = n] \le -\epsilon,$$

for any $n \notin C$. In addition, it is easily verified that $\mathbb{E}[V(N(t+1))|N(t) = n] < \infty$ for any $n \in C$.

Inspection of the Foster-Lyapunov drift criteria [62] then shows that the Markov chain N(t) is positive recurrent, so the system is stable.

REMARK 2.1. If B_k has finite support, i.e., $B_k^{\max} = \sup\{b : \mathbb{P}(B_k = b) > 0\} < \infty$, then the above proof may be considerably simplified by taking $L_k = B_k^{\max}$ and dropping all the terms involving N_l^k , $l \ge B_k^{\max} + 1$.

THEOREM 2.2. For any $\rho < 1$, there exists a scheduling strategy that achieves stability in case of gradual traffic.

PROOF. We introduce several constants that will be used. Let $\epsilon = \frac{1}{2}(1-\rho)/(K+2) > 0$ and define θ_k , L_k and Z_k as in the proof of Theorem 2.1. In addition, define $\varphi_k = \alpha_k \mathbb{E}[D_k]$, $\varphi = \sum_{k=1}^{K} \varphi_k$, $\delta = \epsilon/\varphi$, and $\zeta = (1-\epsilon)/\delta$. Finally, let $M_k = \min\{m : \sum_{j=m+1}^{\infty} j \mathbb{P}(D_k = j) \le \varphi_k/(\alpha_k \zeta)\}$, and observe that $M_k < \infty$ since $\mathbb{E}[D_k] < \infty$.

We consider a scheduling strategy with the following property: it serves an inactive class-*k* flow that either (i) has a feasible transmission rate Z_k or higher; or (ii) has a residual size greater than L_k and a positive feasible transmission rate, whenever possible. Recall that in time slot *t* a flow of class k(t) with a residual size of l(t) bits is served at rate r(t), with the convention that k(t) = l(t) = r(t) = 0 in case no flow gets scheduled in time slot *t* at all.

In order to describe the evolution of the system over time, we denote $N_l^k(t)$ representing the number of inactive class-*k* flows in the system at the beginning of slot *t* with a residual size of *l* bits, and $Q_{lm}^k(t)$ the number of class-*k* flows in the system at time *t* with a residual activity period of length *m* and a total size of *l* bits. The system state is then captured by the vectors of flows $N(t) = (N_1(t), \dots, N_K(t))$, with $N_k(t) = (N_1^k(t), Q_1^k(t), N_2^k(t), Q_2^k(t), \dots)$ and $Q_l^k(t) = (Q_{l1}^k(t), Q_{l2}^k(t), \dots)$.

Observe that

$$N_{l}^{k}(t+1) = N_{l}^{k}(t) + Q_{l1}^{k}(t) - \mathbb{1}_{\{k(t)=k,l(t)=l\}} + \mathbb{1}_{\{k(t)=k,l(t)=l+r(t)\}},$$

and

$$Q_{lm}^{k}(t+1) = Q_{lm+1}^{k}(t) + A_{lm}^{k}(t)$$

with $A_{lm}^k(t)$ denoting the number of class-*k* flows arriving at time *t* with an activity period of length *m* and a size of exactly *l* bits. It is easily verified that the process N(t) is a Markov chain.

Define the Lyapunov function:

$$V(n) = \sum_{k=1}^{K} \left(\delta \sum_{j=1}^{M_{k}} j q_{*j}^{k} + \delta \varsigma \sum_{j=M_{k}+1}^{\infty} j q_{*j}^{k} + \sum_{i=1}^{L_{k}} (q_{i*}^{k} + n_{i}^{k}) \left[\frac{i}{Z_{k}} \right] + \theta_{k} \sum_{i=L_{k}+1}^{\infty} i (q_{i*}^{k} + n_{i}^{k}) \right),$$

INSTABILITY OF MAXWEIGHT SCHEDULING

with $n = (n_1, \ldots, n_K)$, $n_k = (n_1^k, q_1^k, n_2^k, q_2^k, \ldots)$, $q_l^k = (q_{l_1}^k, q_{l_2}^k, \ldots)$, $q_{l_*}^k = \sum_{j=1}^{\infty} q_{l_j}^k$, $q_{*m}^k = \sum_{i=1}^{\infty} q_{im}^k$. The above function provides a measure for the total workload and weighted aggregate residual lifetime of all the flows present in the system.

Note that

$$V(N(t+1)) - V(N(t)) = \sum_{k=1}^{K} I_k(t) + \delta \sum_{k=1}^{K} J_k(t) - \delta \sum_{k=1}^{K} E_k(t) - D(t), \quad (2.6)$$

with

$$I_k(t) = \left(\sum_{i=1}^{L_k} A_{i*}^k(t) \left\lceil \frac{i}{Z_k} \right\rceil + \theta_k \sum_{i=L_k+1}^{\infty} i A_{i*}^k(t) \right),$$

reflecting the increase in the workload due to the arrival of class-*k* flows,

$$J_k(t) = \sum_{j=1}^{M_k} j A_{*j}^k(t) + \zeta \sum_{j=M_k+1}^{\infty} j A_{*j}^k(t)$$

with $A_{i*}^k(t) = \sum_{j=1}^{\infty} A_{ij}^k(t)$, $A_{*j}^k(t) = \sum_{i=1}^{\infty} A_{ij}^k(t)$, representing the increase in the aggregate residual lifetime due to the arrival of class-*k* flows. Moreover,

$$D(t) = \left\lceil \frac{l(t)}{Z_{k(t)}} \right\rceil \mathbb{1}_{\{1 \le l(t) \le L_{k(t)}\}} + \theta_{k(t)} l(t) \mathbb{1}_{\{l(t) > L_{k(t)}\}} - \left\lceil \frac{l(t) - r(t)}{Z_{k(t)}} \right\rceil \mathbb{1}_{\{1 \le l(t) - r(t) \le L_{k(t)}\}} - \theta_{k(t)} (l(t) - r(t)) \mathbb{1}_{\{l(t) - r(t) > L_{k(t)}\}}$$

captures the decrease in the workload due to the service of inactive flows, and

$$E_k(t) = \sum_{j=1}^{M_k} Q_{*j}^k(t) + \zeta \sum_{j=M_k+1}^{\infty} Q_{*j}^k(t)$$

corresponds to the decrease in the aggregate residual lifetime due to the aging of active class-k flows.

Conditioning the drift (2.6) on the number of flows present,

$$\mathbb{E}[V(N(t+1)) - V(N(t)) | N(t) = n]$$

= $\sum_{k=1}^{K} \mathbb{E}[I_k(t)] + \delta \sum_{k=1}^{K} \mathbb{E}[J_k(t)] - \delta \sum_{k=1}^{K} \mathbb{E}[E_k(t) | N(t) = n] - \mathbb{E}[D(t) | N(t) = n].$ (2.7)

Define

$$C = \{n \mid \sum_{k=1}^{K} n_k < N_{\epsilon,\eta} \text{ and } \sum_{k=1}^{K} s_k = 0\}.$$

It may be shown that

$$\mathbb{E}[I_k(t)] \le \rho_k + 2\epsilon, \tag{2.8}$$

$$\mathbb{E}[J_k(t)] \le 2\varphi_k,\tag{2.9}$$

 $\mathbb{E}[E_k(t)|N(t) = n] \ge \zeta, \quad n \notin C, \tag{2.10}$

$$\mathbb{E}[D(t) \mid N(t) = n] \ge 1 - \epsilon \quad n \notin C.$$
(2.11)

2.3 INSTABILITY OF MAXWEIGHT SCHEDULING

Equations (2.8)-(2.11) are derived in Lemma 2.2, which is presented and proven in Appendix 2.A.

27

Define the set

$$\hat{C} = \{n \mid \sum_{k=1}^{K} n_k \le N_{\epsilon,\eta} \text{ and } \sum_{k=1}^{K} s_k = 0 \text{ and } \sum_{k=1}^{K} q_k \le \varsigma \text{ and } \sum_{k=1}^{K} s'_k = 0\}.$$

Suppose $n \notin \hat{C}$. Then either $\sum_{k=1}^{K} q_k > \zeta$ or $\sum_{k=1}^{K} s'_k \ge 1$ or $n \notin C$. If $n \notin C$, then the conditional drift is bounded from above by

$$\rho + 2K\epsilon + 2\delta\varphi - 1 + \epsilon = \rho + (2K + 3)\epsilon - 1 = -\epsilon$$

If $\sum_{k=1}^{K} q_k > \zeta$ or $\sum_{k=1}^{K} s'_k \ge 1$, then the conditional drift is bounded from above by

$$\rho + 2K\epsilon + 2\delta\varphi - \delta\varsigma = \rho + (2K+3)\epsilon - 1 = -\epsilon.$$

Combining Equations (2.7)-(2.11) we obtain

$$\mathbb{E}[V(N(t+1)) - V(N(t)) \mid N(t) = n] \le -\epsilon,$$

for any $n \notin \hat{C}$. In addition, it is easily verified that $\mathbb{E}[V(N(t+1))|N(t) = n] < \infty$ for any $n \in \hat{C}$.

Inspection of the Foster-Lyapunov drift criteria [62] then shows that the Markov chain N(t) is positive recurrent, so the system is stable.

2.3 Instability of MaxWeight scheduling

In this section we establish that MaxWeight scheduling may fail to provide maximum stability. Specifically, we analyze two model instances where the sufficient condition stated in the previous section is satisfied, yet the MaxWeight strategy does not keep the system stable. For the sake of tractability, we focus on relatively simple models with instantaneous traffic and just a single class of flows. In the next section we present extensive simulation results to demonstrate that the instability may also occur in more complex scenarios with gradual traffic that do not lend themselves easily to an analytical treatment.

EXAMPLE 2.1. In this example we consider a single class of flows, and for convenience of notation we omit the subscript indicating the class. Otherwise the notation is identical to that used in Section 2.2. Flows start according to a Bernoulli process, i.e., in each time slot either a flow starts with probability α or no flow starts with probability $1 - \alpha$, independent from slot to slot. The service requirement of each flow is a constant B = 2D + 1 for some integer $D \ge 1$. The feasible transmission rate of a flow is either D + 1 with probability p or 2D + 1 with probability 1 - p, $0 , so <math>R^{\max} = 2D + 1$. The feasible transmission rates are independent across time and among different flows.

It is readily seen that $\tau = 1$ and so $\rho = \alpha$. Theorem 2.1 states that $\rho = \alpha < 1$ is a sufficient condition for stability to be achievable. We now show that the MaxWeight scheduling strategy fails to achieve stability for $\rho = \alpha > 1/(1 + p)$. The reason for the
INSTABILITY OF MAXWEIGHT SCHEDULING

potential instability may be explained as follows. When a flow starts, the MaxWeight strategy will immediately serve it in the next slot, regardless of whether it has feasible rate D + 1 or 2D + 1. To see that, observe that older flows present in the system will necessarily be of size D, and have no chance to be selected in competition with a new flow of size 2D + 1. In case the new flow has feasible rate D + 1, it will require an additional slot at some later point for the service to be completed. In other words, the MaxWeight strategy 'wastes' a second slot on the service of flows whose initial feasible rate is D + 1, whereas a single slot would suffice under a more cautious strategy. More specifically, since the expected number of slots required per flow is 1 + p, it follows that $\alpha > 1/(1 + p)$ precludes stability.

REMARK 2.2. We can extend the example of instability to a slightly more general setting. Consider, as in the situation described above, a system with a single class of flows. Flows start according to a Bernoulli process, i.e., in each time slot either a flow starts with probability α or no flow starts with probability $1 - \alpha$, independent from slot to slot. The service requirement of each flow is a constant *B*. In addition to R^{\max} , we also introduce $R^{\min} = \min\{i : \mathbb{P}(R = i) > 0\}$. Assume now that feasible service rates are such that

$$(B - R^{\min}) \cdot R^{\max} < B \cdot R^{\min}.$$

It is easy to see that this condition implies that a flow entering the system will immediately get scheduled. Hence, the average number of slots required for the service of an arbitrary flow is bounded from below by

$$1 + \sum_{i=1}^{R^{\max}} \left\lceil \frac{B-i}{R^{\max}} \right\rceil \mathbb{P}(R=i).$$
(2.12)

Thus, stability is precluded if

$$\alpha\left(1+\sum_{i=1}^{R^{\max}}\left\lceil\frac{B-i}{R^{\max}}\right\rceil\mathbb{P}(R=i)\right)>1.$$

Note that the quantity in (2.12) is strictly smaller than $\lceil B/R^{\max} \rceil$, provided that $R^{\min} < R^{\max}$.

EXAMPLE 2.2. We discuss a second scenario where the MaxWeight strategy fails to achieve maximum stability. As before, flows start according to a Bernoulli process, i.e., in each time slot either a flow starts with probability α or no flow starts with probability $1 - \alpha$, independent from slot to slot. The service requirement of each flow is a constant *B*. For convenience, we assume B = 8D for some integer $D \ge 1$. The feasible transmission rate of a flow is either 1 with probability *p* or 2 with probability 1 - p, $0 . The feasible transmission rates are independent across time and among different flows. In this case, Theorem 2.1 states that stability can be achieved as long as <math>\rho = 4\alpha D < 1$.

Let $N_i(t)$ denote the number of flows of size *i* at time *t*. It may be shown that for $\rho \le 1$, the process $(N_{3D+1}(t), N_{3D+2}(t), \ldots, N_B(t))$ of flows of size 3D + 1 or greater is stable. This makes sense since large flows receive priority, and the onset of instability manifests itself in the growth of the number of small flows. It then follows that the system spends a non-negligible fraction of time in states where all flows of size 3D + 1

2.4	NUMERICAL	EXPERIMENTS
-----	-----------	-------------

or greater have rate 1 and there is at least one flow of size greater than 6D + 1. In these states, the MaxWeight strategy will serve a flow at rate 1. Similar to the previous scenario, this means that the fraction of time that transmission rate 1 is used, does not approach 0 as $\rho \uparrow 1$, and instability follows.

2.4 Numerical experiments

In this section we present simulation results that confirm the instability of MaxWeight scheduling, as well as clarify the nature of the instability. All simulations consist of a single run of 10^5 time slots. In each slot, a new flow starts with probability α .

The first scenario we consider is Scenario II from Section 2.3, with D = 2. Figure 2.1 shows the number of bits in the system, plotted for various values of α . Although the condition $\alpha < 1$ ensures the existence of a stable scheduling strategy in this scenario, it is easily seen that this is not sufficient for the MaxWeight policy to achieve stability.



Figure 2.1: The number of bits in the system plotted against time under MaxWeight scheduling, for various values of α .

From this point on, we consider gradual traffic. For the duration of the activity period of a flow, a single bit enters in each slot. The length of this period is geometrically distributed with parameter p. In Figure 2.2, three two-class scenarios are presented. Flows belong to either of the classes with equal probability, and the transmission rates are geometrically distributed with parameter q. Hence, $R^{\max} = \infty$, and the necessary stability condition found in Proposition 2.1 simplifies to $\alpha < 1$. Besides the sample path for MaxWeight scheduling, we also plot the behavior of MaxRate scheduling, a somewhat simpler version of the algorithm used in Theorem 2.1 and 2.2, in which the flow with the highest rate is scheduled. In each of these figures, MaxRate scheduling provides stability, whereas MaxWeight scheduling fails to do so. Note that although the MaxWeight scheduling policy is unstable in the cases presented, it is still possible for particular classes of flows to be stable. This is in contrast to MaxWeight scheduling in the static scenario.

Figure 2.3 displays the number of bits over time in a single-class scenario when the transmission rates can assume only two possible values.

INSTABILITY OF MAXWEIGHT SCHEDULING



Figure 2.2: The number of bits in the system of both classes plotted against time for various parameters.

Figure 2.4 contains a similar scenario, but with the transmission rates geometrically distributed with parameter q, so $R^{\max} = \infty$. This figure again demonstrates that MaxWeight fails to provide maximum stability.

2.5 Concluding remarks

We studied the performance of MaxWeight scheduling in a setting where flow dynamics are taken into consideration. We determined an explicit necessary condition for stability, and devised a simple policy to show that this condition in fact is also (nearly) sufficient for stability. Two illustrative examples were provided of scenarios where MaxWeight scheduling fails to attain stability under this condition. The analytical results are supported and complemented by simulation experiments for more involved scenarios. The simulations compare the MaxWeight scheduling algorithm to MaxRate scheduling, and confirm the instability of MaxWeight scheduling.

It is crucial to observe that the rate variations play a critical role in the instability

2.5 CONCLUDING REMARKS



31

Figure 2.3: The number of bits in the system plotted against time for various parameters.

results of this chapter. Intuitively speaking, MaxWeight policies tend to favor flows with large backlogs, even when their service rates are not particularly favorable, and thus fail to maximally exploit the rate variations of flows with smaller backlogs. This raises the question whether the rate variations are essential for the instability to occur. In the case of a shared downlink, where only a single flow can be scheduled at a time, the instability cannot occur in the absence of any rate variations, since this system is work-conserving, and any non-idling scheduling strategy will in fact achieve maximum stability. In the next chapter however, we will demonstrate that in a spatial setting, the instability can occur even without rate variations.

It is worth emphasizing that the scheduling strategies considered in the proofs of Theorems 2.1 and 2.2 mainly serve to prove that $\rho < 1$ is sufficient for the existence of a stable strategy, and are therefore specifically designed for that purpose. The strategies may not be ideal for practical purposes as they may not provide particularly good performance, especially at lower loads. They also involve knowledge of various parameter values, which may be hard to obtain and is not used by the MaxWeight policy. (While the latter may be considered 'unfair', observe that in the standard case with a fixed set of flows no amount of additional information can help to achieve

INSTABILITY OF MAXWEIGHT SCHEDULING



Figure 2.4: The number of bits in the system plotted against time for various parameters.

better stability performance than the MaxWeight policy provides.) The fact that for gradual traffic the scheduling strategy assumes prior knowledge of the duration of the activity period further adds to this.

Recently various throughput-optimal schedulers have been proposed that require less information on the network parameters [54, 55, 74]. The schedulers considered in [54, 55] do this by learning over time the distribution of the feasible transmission rate of each flow. Although this does eliminate the need for explicit information on this distribution, the process of learning the rate distributions takes time, which has an adverse effect on the flow delay. In [74] delay-based MaxWeight scheduling is proposed, i.e., replace the queue length in (1.4) by the waiting time of the head of line packet. It has been shown to hold that this policy is throughput-optimal, but only in the regime with a single class of flows.

2.A AUXILIARY RESULTS

Appendix

2.A Auxiliary results

LEMMA 2.1. In the case of instantaneous traffic, and under the policy defined in the proof of Theorem 2.1

33

$$\mathbb{E}[I_k(t)] \le \rho_k + 2\epsilon,$$
$$\mathbb{E}[D(t) \mid N(t) = n] \ge 1 - \epsilon, \quad n \notin C.$$

PROOF. Note that

$$\begin{aligned} \alpha_k \mathbb{E}[[B_k/Z_k]] &\leq \alpha_k \max\{\mathbb{E}[[B_k/R_k^{\max}], \mathbb{E}[[\epsilon B_k/(\alpha_k \mathbb{E}[B_k])]]\} \\ &\leq \alpha_k \max\{\max\{\mathbb{E}[[B_k/R_k^{\max}]], 1\}\}, 1 + \epsilon/\alpha_k\} \leq \max\{\rho_k, \alpha_k + \epsilon\} \\ &\leq \rho_k + \epsilon. \end{aligned}$$

We first derive an upper bound for $\mathbb{E}[I_k(t)]$. By rewriting (2.1) we obtain

$$\mathbb{E}[I_{k}(t)] = \sum_{i=1}^{L_{k}} \left\lceil \frac{i}{Z_{k}} \right\rceil \mathbb{E}[A_{i}^{k}(t)] + \theta_{k} \sum_{i=L_{k}+1}^{\infty} i\mathbb{E}[A_{i}^{k}(t)]$$

$$= \alpha_{k} \left(\sum_{i=1}^{L_{k}} \left\lceil \frac{i}{Z_{k}} \right\rceil \mathbb{P}(B_{k}=i) + \theta_{k} \sum_{i=L_{k}+1}^{\infty} i\mathbb{P}(B_{k}=i) \right)$$

$$= \alpha_{k} \left(\mathbb{E}\left[\left\lceil \frac{B_{k}}{Z_{k}} \right\rceil \right] + \theta_{k} \sum_{i=L_{k}+1}^{\infty} (i - \left\lceil \frac{i}{Z_{k}} \right\rceil / \theta_{k}) \mathbb{P}(B_{k}=i) \right) \le \rho_{k} + 2\epsilon.$$
(2.13)

From (2.2) it can be seen that

$$D(t) \geq \mathbb{1}_{\{1 \leq l(t) \leq L_{k(t)}, r(t) \geq Z_{k(t)}\}} + \theta_{k(t)} \mathbb{1}_{\{l(t) > L_{k(t)}, r(t) > 0\}}$$

= $\mathbb{1}_{\{1 \leq l(t) \leq L_{k(t)}, r(t) \geq Z_{k(t)}\}} + [\mathbb{P}(R_{k(t)} > 0)]^{-1} \mathbb{1}_{\{l(t) > L_{k(t)}, r(t) > 0\}}.$ (2.14)

Now, let $E^{small}(t)$ be the event that there is at least one class-*k* flow in time slot *t* of residual size no larger than L_k with feasible transmission rate Z_k or higher. Let $E^{large}(t)$ be the event that there is at least one class-*k* flow in time slot *t* of residual size $L_k + 1$ or larger with a non-zero feasible transmission rate.

Note that

$$\mathbb{1}_{\{l(t)>L_{k(t)},r(t)>0\}} = \mathbb{1}_{\{l(t)\geq L_{k(t)}+1,r(t)>0,E^{small}(t)\}} + \mathbb{1}_{\{l(t)>L_{k(t)},r(t)>0,\bar{E}^{small}(t)\}}
= \mathbb{1}_{\{l(t)>L_{k(t)},r(t)>0,E^{small}(t)\}} + \mathbb{1}_{\{E^{large}(t),\bar{E}^{small}(t)\}}.$$
(2.15)

Further observe

$$\mathbb{1}_{\{l(t) \le L_{k(t)}, r(t) \ge Z_{k(t)}\}} + \mathbb{1}_{\{l(t) > L_{k(t)}, E^{small}(t)\}} = \mathbb{1}_{\{E^{small}(t)\}} = 1 - \mathbb{1}_{\{\bar{E}^{small}(t)\}}.$$
(2.16)

Combining (2.14)-(2.16) we deduce that

$$D(t) \ge 1 - \mathbb{1}_{\{\bar{E}small(t)\}} + [\mathbb{P}(R_{k(t)} > 0)]^{-1} \mathbb{1}_{\{E^{large}(t), \bar{E}small(t)\}}.$$

INSTABILITY OF MAXWEIGHT SCHEDULING

Thus,

$$\mathbb{E}[D(t)|N(t) = n] = 1 - \mathbb{P}(\bar{E}^{small}(t)|N(t) = n) + [\mathbb{P}(R_{k(t)} > 0)]^{-1}\mathbb{P}(E^{large}(t), \bar{E}^{small}(t)|N(t) = n) = 1 - \mathbb{P}(\bar{E}^{small}(t)|N(t) = n)(1 - [\mathbb{P}(R_{k(t)} > 0)]^{-1}\mathbb{P}(E^{large}(t)|N(t) = n)).$$
(2.17)

Let $s_k = \sum_{i=L_k}^{\infty} n_i^k$. If $\sum_{k=1}^K s_k > 0$, then $\mathbb{P}(E^{large}(t)|N(t) = n) \ge \min_{k=1,\dots,K} \mathbb{P}(R_k > 0)$, so that $[\mathbb{P}(R_{k(t)} > 0)]^{-1}\mathbb{P}(E^{large}(t)|N(t) = n) \ge 1$. If $\sum_{k=1}^K n_k \ge N_{\epsilon,\eta}$, then $\mathbb{P}(\bar{E}^{small}(t)|N(t) = n) \le (1 - \eta)^{N_{\epsilon,\eta}} \le \epsilon$. Then we obtain from (2.17) that

$$\mathbb{E}[D(t) \mid N(t) = n] \ge 1 - \epsilon \tag{2.18}$$

for any $n \notin C$.

LEMMA 2.2. In the case of gradual traffic, and under the policy defined in the proof of *Theorem 2.2*

$$\mathbb{E}[I_k(t)] \le \rho_k + 2\epsilon,$$
$$\mathbb{E}[J_k(t)] \le 2\varphi_k,$$
$$\mathbb{E}[E_k(t)|N(t) = n] \ge \zeta, \quad n \notin C,$$
$$\mathbb{E}[D(t)|N(t) = n] \ge 1 - \epsilon \quad n \notin C.$$

PROOF. As the arrival process of new flows is the same for both instantaneous traffic and gradual traffic, we conclude from (2.4) that

$$\mathbb{E}[I_k(t)] \le \rho_k + 2\epsilon.$$

Next we establish an upper bound for $\mathbb{E}[J_k(t)]$:

$$\mathbb{E}[J_k(t)] = \sum_{j=1}^{M_k} j\mathbb{E}[A_{*j}^k(t)] + \zeta \sum_{j=M_k+1}^{\infty} j\mathbb{E}[A_{*j}^k(t)]$$
$$= \alpha_k \Big(\sum_{j=1}^{M_k} j\mathbb{P}(D_k = j) + \zeta \sum_{j=M_k+1}^{\infty} j\mathbb{P}(D_k = j)\Big)$$
$$= \alpha_k \Big(\mathbb{E}[D_k] + \zeta \sum_{j=M_k+1}^{\infty} j(1 - 1/\zeta)\mathbb{P}(D_k = j)\Big) \le 2\varphi_k$$

We proceed with a lower bound for $\mathbb{E}[E_k(t)|N(t) = n]$:

$$\mathbb{E}[E_k(t)|N(t) = n] = \sum_{j=1}^{M_k} q_{*,j}^k + \zeta \sum_{j=M_k+1}^{\infty} q_{*,j}^k$$

Thus $\mathbb{E}[E_k(t)|N(t) = n] \ge \varsigma$ whenever $q_k = \sum_{j=1}^{M_k} q_{*j}^k \ge \varsigma$ or $s'_k = \sum_{j=M_k+1}^{\infty} q_{*j}^k \ge 1$. We turn to a lower bound for $\mathbb{E}[D(t) \mid N(t) = n]$. Recall

$$C = \{n \mid \sum_{k=1}^{K} n_k < N_{\epsilon,\eta} \text{ and } \sum_{k=1}^{K} s_k = 0\}.$$

2.A AUXILIARY RESULTS

Note that

$$\begin{split} D(t) \geq & \mathbb{1}_{\{1 \leq l(t) \leq L_{k(t)}, r(t) \geq Z_{k(t)}\}} + \theta_{k(t)} \mathbb{1}_{\{l(t) > L_{k(t)}, r(t) \geq 1\}} \\ = & \mathbb{1}_{\{1 \leq l(t) \leq L_{k(t)}, r(t) \geq Z_{k(t)}\}} + [\mathbb{P}(R_{k(t)} > 0)]^{-1} \mathbb{1}_{\{l(t) > L_{k(t)}, r(t) > 0\}}. \end{split}$$

Now, let $E^{small}(t)$ be the event that there is at least one class-*k* flow in time slot *t* of size no larger than L_k with feasible transmission rate Z_k or larger. Let $E^{large}(t)$ be the event that there is at least one class-*k* flow in time slot *t* of size $L_k + 1$ or larger with positive feasible transmission rate.

Note that

$$\begin{split} \mathbb{1}_{\{l(t)>L_{k(t)},r(t)>0\}} &= \mathbb{1}_{\{l(t)>L_{k(t)},r(t)>0,E^{small}(t)\}} + \mathbb{1}_{\{l(t)>L_{k(t)},r(t)>0,\bar{E}^{small}(t)\}} \\ &= \mathbb{1}_{\{l(t)>L_{k(t)},r(t)>0,E^{small}(t)\}} + \mathbb{1}_{\{E^{large}(t),\bar{E}^{small}(t)\}}. \end{split}$$

Further observe

$$\mathbb{1}_{\{l(t) \le L_{k(t)}, r(t) \ge Z_{k(t)}\}} + \mathbb{1}_{\{l(t) > L_{k(t)}, E^{small}(t)\}} = \mathbb{1}_{\{E^{small}(t)\}} = 1 - \mathbb{1}_{\{\bar{E}^{small}(t)\}}.$$

We deduce that

$$D(t) \ge 1 - \mathbb{1}_{\{\bar{E}^{small}(t)\}} + [\mathbb{P}(R_{k(t)} > 0)]^{-1} \mathbb{1}_{\{E^{large}(t), \bar{E}^{small}(t)\}}.$$

Thus,

$$\begin{split} \mathbb{E}[D(t)|N(t) = n] = &1 - \mathbb{P}(\bar{E}^{small}(t)|N(t) = n) \\ &+ [\mathbb{P}(R_{k(t)} > 0)]^{-1}\mathbb{P}(E^{large}(t), \bar{E}^{small}(t)|N(t) = n) \\ &= &1 - \mathbb{P}(\bar{E}^{small}(t)|N(t) = n) \\ &\cdot (1 - [\mathbb{P}(R_{k(t)} > 0)]^{-1}\mathbb{P}(E^{large}(t)|N(t) = n)). \end{split}$$

If $\sum_{k=1}^{K} s_k > 0$, then $\mathbb{P}(E^{large}(t)|N(t) = n) \ge \min_{k=1,\dots,K} \mathbb{P}(R_k > 0)$, so that $[\mathbb{P}(R_{k(t)} > 0)]^{-1}\mathbb{P}(E^{large}(t)|N(t) = n) \ge 1$. If $\sum_{k=1}^{K} n_k \ge N_{\epsilon,\eta}$, then $\mathbb{P}(\bar{E}^{small}(t)|N(t) = n) \le (1 - \eta)^{N_{\epsilon,\eta}} \le \epsilon$.

We then obtain

$$\mathbb{E}[D(t) \mid N(t) = n] \geq 1 - \epsilon$$

for any $n \notin C$.

~	5	
υ	J	

36 INSTABILITY OF MAXWEIGHT SCHEDULING

3

SPATIAL INEFFICIENCY OF MAXWEIGHT SCHEDULING

In the previous chapter we have seen that the MaxWeight scheduling algorithm may fail to achieve maximum stability in a setting with flow-level dynamics. The instability examples in that chapter all consider a single-downlink wireless channel with time-varying transmission rates. The more challenging problem, however, arises in networks where certain subsets of the links can be activated simultaneously subject to interference constraints. In the present chapter we show that MaxWeight scheduling policies may fail to provide maximum stability in such scenarios as well, even in the absence of any rate variations. We show that the potential instability effects can be countered by implementing a region-based version of MaxWeight scheduling.

This chapter is organized as follows. In Section 3.1 we provide a detailed model description, and in Section 3.2 we demonstrate the potential instability of MaxWeight scheduling through several examples. In Section 3.3 we examine the performance of region-based scheduling in two-dimensional networks with an arbitrary spatial traffic density. Section 3.4 offers some concluding remarks.

SPATIAL INEFFICIENCY OF MAXWEIGHT SCHEDULING

3.1 Model description

We consider a time-slotted wireless system on some space *S*. Traffic consists of finitesized flows that enter the system at random, and leave once fully served. Each arriving flow is associated with a certain location in *S* and a finite size, as will be further described for specific model instances later. In each time slot, a centralized scheduler selects a subset of flows for transmission. In the present chapter we assume for simplicity that the total size of a flow is known upon arrival, and no further traffic of that flow will arrive. Most of the results can be extended to a setting with gradual traffic, where traffic of a flow arrives over time.

A subset of points in *S* is said to be feasible if flows in these locations can be scheduled simultaneously. The function $F(\cdot)$ indicates whether or not a subset of points is feasible, i.e., given *n* flows with distinct locations $X_1, \ldots, X_n \in S$, $F(\{X_1, \ldots, X_n\})$ equals 1 if these flows can be scheduled simultaneously and is 0 otherwise. Flows in the same location can never be scheduled simultaneously. A prototypical scenario would be that $F(\{X_1, \ldots, X_n\}) = 1$ if and only if $||X_i - X_j|| \ge d$ for all $i \ne j$, which corresponds to a reuse distance *d*, and is similar in spirit as the protocol model. However, the feasibility function could also be based on the physical model for example.

In each time slot a certain subset of flows gets selected for service, as governed by the applicable scheduling strategy, subject to the feasibility constraints. Each time a flow gets scheduled, its residual size is reduced by 1, and a flow leaves the system once it has been served to completion i.e., its size reaches 0. The subset of flows selected by the scheduling strategy may depend on the locations $X_i(t)$ and residual sizes $Q_i(t)$, $i \in I(t)$, with I(t) indexing the flows present in time slot t. In particular, the MaxWeight scheduling strategy selects a feasible subset of flows $J^*(t) \subseteq I(t)$, $F(J^*(t)) = 1$, of maximum aggregate residual size, i.e.,

$$\sum_{j \in J^*(t)} Q_j(t) = \max_{J \subseteq I(t), F(J) = 1} \sum_{j \in J} Q_j(t).$$
(3.1)

The main reason for assuming unit transmission rates is to stress the fact that the instability phenomena demonstrated in later sections result from persistent spatial inefficiency rather than rate heterogeneity. Possible rate heterogeneity induces priorities among flows, which may exacerbate the spatial inefficiency and render the system even more prone to potential instability effects.

3.2 Instability of MaxWeight Scheduling

In this section we present several illustrative examples where the MaxWeight scheduling strategy fails to achieve maximum stability.

EXAMPLE 3.1. We first consider a network with three regions as shown in Figure 3.1. Transmissions in region 2 interfere with transmissions in both region 1 and region 3, and transmissions in regions 1 and 3 do not interfere with each other. Flows arrive at region *i* at a rate λ_i (per time slot) and have initial size B_i . Denote by $\rho_i = \lambda_i \mathbb{E}[B_i]$ the traffic intensity at region *i*. We assume that

$$\rho_1 + \rho_2 < 1 \text{ and } \rho_3 + \rho_2 < 1,$$

3.2 INSTABILITY OF MAXWEIGHT SCHEDULING

or equivalently,

$$\rho_2 < 1 - \max\{\rho_1, \rho_3\}. \tag{3.2}$$

It is easily seen that the latter condition is necessary for stability to be achievable, and in fact also sufficient under mild independence assumptions. A scheduling strategy that can stabilize the network when (3.2) holds is as follows. At each time slot, schedule a flow in region 2 with probability $\rho_2 + \epsilon$, or schedule a flow in both regions 1 and 3, with probability $\max\{\rho_1, \rho_3\} + \epsilon$, where $\epsilon = \frac{1-\rho_2 - \max\{\rho_1, \rho_3\}}{2}$.

Now suppose $B_2 \equiv 1$, and recall that the MaxWeight scheduling strategy as defined in the general network model of the previous section selects a set of *flows* with maximum aggregate residual size. Thus the MaxWeight strategy never schedules a flow in region 2 as long as a flow with a residual size of 2 or larger is present in region 1 or region 3. Hence the scheduling of flows of residual size 2 or larger in region 1 and region 3 is independent from each other. Also, the fraction of time that a flow of residual size 2 or larger gets scheduled in region *i*, is $\lambda_i(\mathbb{E}[B_i] - 1) = \rho_i - \lambda_i$. It follows that the fraction of time that a flow in region 2 gets scheduled, is bounded from above by

$$(1 - \rho_1 + \lambda_1)(1 - \rho_3 + \lambda_3).$$

Thus a necessary condition for MaxWeight scheduling to achieve stability is $\rho_2 \leq$ $(1 - \rho_1 + \lambda_1)(1 - \rho_3 + \lambda_3)$. When the λ_i 's (i = 1, 3) are small and the $\mathbb{E}[B_i]$'s (i = 1, 3)are large, the latter condition 'approaches' $\rho_2 \leq (1 - \rho_1)(1 - \rho_3)$, which is a more stringent inequality than the sufficient condition (3.2).



region 3



In Example 3.1, the stabilizing strategy either schedules both region 1 and region 3, or schedules region 2. The MaxWeight policy, however, tends to serve flows with large backlogs, so flows in regions 1 and 3 are served with priority when their residual sizes are greater than or equal to 2. Consequently, the MaxWeight policy schedules a flow

SPATIAL INEFFICIENCY OF MAXWEIGHT SCHEDULING

in region 1 (region 3) even when region 3 (region 1) is empty, which leads to inefficient spatial reuse. Thus, the MaxWeight policy fails to achieve maximum stability.

Example 3.1 illustrates the spatial inefficiency of MaxWeight scheduling by carefully constructing the regions where flows arrive. Next we present a further example where we consider a one-dimensional space (a ring) with uniformly distributed arrival locations. We assume that all flows are of the same size, and show that even in this *uniform* traffic scenario, MaxWeight scheduling fails to achieve throughput optimality.

EXAMPLE 3.2. Let $N \ge 1$ and consider a ring with unit circumference and reuse distance d = 2(N + 1)/((2N + 3)(3N + 2)), partitioned into (2N + 3)(3N + 2) intervals of equal size, see Figure 3.2. In each time slot, either exactly (2N + 3) flows arrive with probability a, each of size B = 2, at locations uniformly distributed in the intervals M + j(3N + 2), j = 1, 2, ..., (2N + 3), where M is uniformly distributed on 1, 2, ..., 3N + 2, or no flows arrive at all with probability 1 - a.



Figure 3.2: A ring with unit circumference, reuse distance d = 4/25, partitioned into 25 intervals of equal size (N = 1).

Consider a strategy that generates a random variable *L* uniformly distributed on 1, 2, ..., 2*N*+3, and then selects an arbitrary flow for service from each of the intervals *L* + *i*(2*N* + 3), *i* = 0, 1, ..., 3*N* + 1, if available. Note that the strategy respects the reuse distance, and achieves stability as long as the aggregate traffic intensity in each interval, 2a/(3N + 2), is less than the fraction of time slots that each interval gets selected for service, 1/(2N + 3), or equivalently, if a < a(N) = (3N + 2)/(4N + 6). Note that $a(N) \rightarrow 3/4$ as $N \rightarrow \infty$. Also, the maximum size of a feasible subset of points is $M(N) = \lfloor \frac{(2N+3)(3N+2)}{2(N+1)} \rfloor$, and the total traffic intensity equals $\rho = 2a(2N + 3)$, so the necessary condition $\rho < M$ for stability takes the form

$$a < b(N) = \frac{1}{2(2N+3)} \left\lfloor \frac{(2N+3)(3N+2)}{2(N+1)} \right\rfloor.$$

Observe that $b(N) \rightarrow 3/4$ as $N \rightarrow \infty$, and thus the above-described strategy in fact achieves maximum stability for large values of *N*.

It is easily verified that in each time slot with arriving flows, the MaxWeight strategy selects all 2N + 3 of them for service, while in a time slot without any arrivals, it can

3.3 STABILITY OF REGION-BASED SCHEDULING

serve at most 3N + 2 traffic units, so the expected total number of traffic units served per time slot is bounded from above by a(2N + 3) + (1 - a)(3N + 2). As a necessary condition in order for the MaxWeight strategy to be stable, the latter number must be larger than the total traffic intensity 2a(2N + 3), which entails $a < a^{MW}(N) = (3N + 2)/(5N + 5)$. Note that $a^{MW}(N) \le a(N)$, with strict inequality for all $N \ge 2$, and that $a^{MW}(N) \to 3/5$ as $N \to \infty$.

41

We conclude that for $a \in (a^{MW}(N), a(N))$, the MaxWeight strategy fails to achieve stability, although there exists a strategy that does provide stability. For large values of N, the MaxWeight strategy is only able to sustain at most a fraction 4/5 of the maximum throughput.

In the above example MaxWeight scheduling always selected newly arrived flows for transmission, even when it could have chosen a subset that allowed for better spatial reuse. This persistent inefficiency then leads to instability. As we see below, this behavior occurs for more general traffic patterns as well.

The locations of arriving flows in Example 3.2 are uniformly distributed, but highly correlated. When the flow locations are independent, the behavior is more complex, and (in)stability is more difficult to establish. We therefore proceed with a simulation experiment where we assume that the locations of arriving flows are independent. As we show in the next example, the MaxWeight strategy again fails to achieve throughput optimality.

EXAMPLE 3.3. Consider a ring network where the total number of arriving flows is geometrically distributed with parameter p = 0.45 and mean $\alpha = (1 - p)/p \approx 11/9$, so $\rho = \alpha \mathbb{E}[B] \approx 22/9$. We assume that the locations of the flows are independent and uniformly distributed along the ring. The reuse distance is d = 0.3, and hence the maximal number of flows that can be scheduled simultaneously equals M = 3.

We compare the performance of MaxWeight scheduling with that of a randomized interval-based scheduling strategy. We divide the ring into 42 intervals of length 1/42. We consider 42 schedules $\omega_k = \{k, k + 14, k + 28\}$ (modulo 42), and choose in each time slot one of these schedules uniformly at random.

We simulate the network 1000 slots, for both MaxWeight scheduling and the randomized strategy. Figure 3.3 shows the total number of flows present over time for MaxWeight scheduling (gray) and the randomized strategy (black). Under MaxWeight scheduling the number of flows grows without bound, suggesting instability. In contrast, the number of flows settles around a relatively low level for the randomized strategy.

3.3 Stability of region-based scheduling

In the previous section we demonstrated the spatial inefficiency of MaxWeight scheduling. This raises the question of finding scheduling algorithms that can be used to stabilize spatial networks with flow-level dynamics. For the single-channel case with flow-level dynamics a (impractical) stabilizing policy was presented in Chapter 2. Moreover, it was recently shown that maximum stability can be achieved by scheduling according to the feasible transmission rate [54, 55] or according to the product of feasible transmission rate and the delay [74]. It is not clear whether these policies are throughput-optimal in the spatial setting, or how they may need be modified





Figure 3.3: Evolution of the total number of flows for MaxWeight scheduling (gray) and an interval-based randomized scheduler (black), respectively.

to provide maximum stability. Both schedulers have to find a maximum (weighted) independent set over all flows in each time slot, and since the number of flows is unbounded, so is the scheduling complexity. In this section we present a class of policies that do have bounded complexity.

Consider a two-dimensional network with an arbitrary spatial traffic density. We assume that the space *S* is bounded, and without loss of generality we may suppose that location coordinates are scaled such that *S* is contained in the unit square $[0,1]^2$. The number of arriving flows, their locations, and their sizes are independent and identically distributed across time slots. The location of an arbitrary arriving flow is governed by some spatial measure λ on $[0,1]^2$, with $\lambda(x, y) = 0$ for all $(x, y) \notin S$, i.e., the expected number of arriving flows per time slot in a region $\mathcal{R} \subseteq S$ is $\int_{(x,y)\in\mathcal{R}} \lambda(x, y) \, dx \, dy$. Let the positive random variable *B* represent the size of an arbitrary flow.

3.3.1 Two special cases

The above general network setting includes the three-region network and the ring topology with uniform traffic density discussed in Section 3.2. Before presenting results on the stability of general spatial networks with flow-level dynamics, let us first consider maximal stable policies for these two special cases. For the three-region network, an alternative view of the network is to consider each region as a single node. The three-region network can then be seen as a classic three-node network, where packets belonging to various flows are continuously injected into each node. It is easily verified that in this case the MaxWeight strategy that schedules according to the aggregate backlog at each node is throughput-optimal, see [89].

Now consider the ring network with uniform spatial traffic density α . Taking a similar approach as for the three-region network, instead of scheduling flows, we

3.3 STABILITY OF REGION-BASED SCHEDULING

divide the ring into intervals and schedule these intervals instead. As we show in the following proposition, for the right choice of intervals the ring network can be stabilized using such interval-based scheduling.

PROPOSITION 3.1. Consider a ring with unit circumference, uniform spatial traffic density, and a translation-invariant feasibility function, and denote by M the maximum number of flows that can be scheduled simultaneously. Then for any load $\rho = \alpha \mathbb{E}[B] < M$ there exists an interval-based scheduling strategy that achieves stability.

PROOF. Let $\mathcal{P} = \{X_1, \ldots, X_M\} \subseteq [0, 1]$ denote a feasible maximum-size set and assume that there exist some $\epsilon > 0$ such that for any $\hat{X}_i \in [X_i, X_i + \epsilon)$, the set $\hat{\mathcal{P}} = \{\hat{X}_1, \ldots, \hat{X}_M\}$ is feasible as well. We choose integers K_i , $i = 1, 2, \ldots, M$ and K such that each interval $[X_i, X_i + \epsilon)$ contains the points K_i/K and $(K_i + 1)/K$, $i = 1, 2, \ldots, M$. We partition the ring into K intervals, each of size 1/K. Now consider a cyclic scheduling strategy which in time slot tK + u, $u = 1, \ldots, K$, $t = 0, 1, \ldots$, serves the intervals $[(K_i + u)/K, (K_i + 1 + u)/K]$, $i = 1, \ldots, M$, by selecting an arbitrary flow from each of these intervals, if available. Note that any set of flows thus selected is allowed since the feasibility function is translation-invariant. Also, each interval is allowed to be served a fraction of the time M/K and has aggregate traffic intensity ρ/K . Hence, the strategy achieves stability for any $\rho < M$.

Note that Proposition 3.1 assumes a general interference model, so it includes the model with reuse distance discussed in Examples 3.2 and 3.3 as a special case. Consider the case where the reuse distance is *d*, and assume 1/d is not an integer. Then $M = \lfloor 1/d \rfloor$ and $\epsilon = 1/M - d$. It can be readily verified that given the reuse distance *d*, the necessary condition for ρ to be supportable is $\rho \leq \lfloor 1/d \rfloor$ and that the scheduling algorithm presented in the proof can stabilize any $\rho < \lfloor 1/d \rfloor$.

3.3.2 General networks

The examples in Section 3.3.1 indicate that in the presence of flow-level dynamics we should aggregate over several nearby flows, rather than schedule based on individual flows. This suggests a *region-based* scheduling algorithm where the space is partitioned into a finite number of regions. In each time slot, the algorithm selects a subset of non-interfering regions and then schedules a flow in each selected region, if any.

Naturally, such a partitioning would reduce the flexibility of the scheduler since the region-based feasibility constraints are more stringent than the original constraints. Region-based scheduling is nevertheless useful because, in contrast to the partition-free system, throughput-optimal schedulers are available in this case. For example, since the partitioned system behaves as a network with a finite number of persistent queues, it is well-known that region-based MaxWeight scheduling (i.e., MaxWeight scheduling based on the aggregate backlog of all flows in a region) is throughput-optimal within the class of schedulers that satisfy the more stringent feasibility constraints of the partitioned system. We are interested in how the capacity region of the partitioned system relates to the capacity region under the original reuse constraints. As we will see, this depends on the granularity of the partitioning. Note that region-based MaxWeight scheduling complexity, as the number of regions is fixed.

SPATIAL INEFFICIENCY OF MAXWEIGHT SCHEDULING

We continue to consider a specific form of region-based scheduling referred to as *K*-partition, where the area $[0,1]^2$ is partitioned into K^2 square cells of size $1/K^2$, for some $K \in \mathbb{N}$. The cells are denoted $\mathcal{R}_{k,l} = [(k-1)/K, k/K] \times [(l-1)/K, l/K]$, $k, l = 1, \ldots, K$. The 4-partition is illustrated in Figure 3.4.

$\mathcal{R}_{1,4}$		$\mathcal{R}_{4,4}$
$\mathcal{R}_{1,1}$		$\mathcal{R}_{4,1}$

Figure 3.4: The 4-partition, where the unit square is divided into 16 cells.

Under *K*-partition, a set of cells $\mathcal{R}_{k_1,l_1}, \mathcal{R}_{k_2,l_2}, \ldots, \mathcal{R}_{k_n,l_n}$ is said to be feasible if for any $X_i \in \mathcal{R}_{k_i,l_i}$, $i = 1, \ldots, n$, the set of points $C = \{X_1, X_2, \ldots, X_n\}$ is feasible. For convenience, we henceforth assume that the feasibility function is governed by a reuse distance *d* constraint. Denote by $\Omega(d)$ all feasible sets of points, and let $\Omega(K, d) \subseteq$ $\{0, 1\}^{K^2}$ represent the collection of all feasible subsets of cells. So for $\omega \in \Omega(K, d)$ we have that $\omega_{k,l} = 1$ if $\mathcal{R}_{k,l}$ is contained in the schedule ω , and $\omega_{k,l} = 0$ otherwise. We focus on square regions for convenience, but we expect that similar qualitative results hold for other types of regions.

Under *K*-partition, scheduling is confined to subsets of flows that belong to a feasible subset of cells, which restricts beyond the original reuse constraint and guarantees feasibility. The aggregate arrival rate of flows into the cell $\mathcal{R}_{k,l}$ is given by $\lambda_{k,l} = \int_{(x,y)\in\mathcal{R}_{k,l}} \lambda(x, y) \, dx dy$. The capacity region for such a system is well-known:

$$C(K,d) = \{\lambda : \lambda_{k,l} \mathbb{E}[B] \in \text{conv.hull}(\Omega(K,d))\}.$$

Let C(d) denote the capacity region under the original reuse constraint, then $C(K, d) \subseteq C(d)$ for any $K \ge 1$. As K increases, the granularity of the partitioning becomes finer, and it is intuitive that C(K, d) converges to C(d) in a certain sense. This is formalized in Theorem 3.1, which states our main result, showing that for any arrival density function $\lambda \in C(d)$ (under certain assumptions) there exists a K such that λ is contained in C(K, d).

Before stating Theorem 3.1, we first present the following lemmas.

LEMMA 3.1. Let d > 0 and $K \in \mathbb{N}$, $K \ge 2\sqrt{2}/d$, then

$$C(d) \subseteq C(K, d - 2\sqrt{2}/K).$$

The proof of Lemma 3.1 is presented in Appendix 3.A.1.

Let $\omega \in \Omega(K, d)$, $L \leq K$, and denote by $\omega^{(L)}$ the vector ω restricted to the entries $\omega_{k,l}$, k, l = 1, 2, ..., L. Then the following lemma holds.

3.3 STABILITY OF REGION-BASED SCHEDULING

LEMMA 3.2. Let d > 0, $K \in \mathbb{N}$ and set

$$h = K \left[\frac{K(d - 2\sqrt{2}/K)}{d} \right]^{-1}.$$
 (3.3)

Then $\omega \in \Omega(K, d) \Rightarrow \omega^{(K/h)} \in \Omega(K/h, d)$.

Lemma 3.2 states that if ω is a feasible set of cells under *K*-partition and reuse distance $d - 2\sqrt{2}/K$, then it is a feasible set of cells under (K/h)-partition and reuse distance *d* as well. The proof of the lemma is presented in Appendix 3.A.2.

We are now in a position to state and prove Theorem 3.1. An arrival density function λ is said to be smooth if it is

- uniformly lower bounded, i.e., there exists a $\kappa^{(0)} > 0$ such that $\lambda(x, y) \ge \kappa^{(0)}$ for all $(x, y) \in S$;
- differentiable, with a uniformly upper bounded first-order partial derivative, i.e., there exists a $\kappa^{(1)} < \infty$ such that $\frac{\partial \lambda(x,y)}{\partial x} \le \kappa^{(1)}$ and $\frac{\partial \lambda(x,y)}{\partial y} \le \kappa^{(1)}$ for all $(x, y) \in S$.

THEOREM 3.1. Let λ be a smooth arrival density function such that $(1 + \epsilon)\lambda \in C(d)$ for some $\epsilon > 0$. Then there exists a $K = K(\lambda)$ such that $\lambda \in C(K, d)$.

The proof of Theorem 3.1 is presented in Appendix 3.A.3. The idea behind the proof of Theorem 3.1 is as follows. By Lemma 3.1 we know that for any given arrival density function within the capacity region C(d), the system can be stabilized by a randomized region-based algorithm under *K*-partition and reduced reuse distance $d - 2\sqrt{2}/K$ that selects schedule $\omega \in \Omega(K, d - 2\sqrt{2}/K)$ with a certain probability $\pi(\omega)$. In order to turn this mechanism into a scheduler that is feasible for reuse distance *d*, we scale the entire system by a factor h^{-1} , and by Lemma 3.2 we know that our randomized scheduler is now valid for reuse distance *d*. This is illustrated in Figure 3.5 for the 8-partition. Certain cells in the scaled system are located outside the unit square, and scheduling them does not result in flows being served. However, by choosing *K* sufficiently large we can make this throughput loss arbitrarily small, thus stabilizing the system.



Figure 3.5: Constructing a 6-partition from the original 8-partition.

While we have demonstrated in Theorem 3.1 that the capacity region C(K, d) of the partitioned system approaches C(d) in a certain sense as K increases, it can be

SPATIAL INEFFICIENCY OF MAXWEIGHT SCHEDULING

shown that they never coincide. In particular, we next establish a negative result which states that for any given *K*-partition, one can construct an arrival density function $\tilde{\lambda}$ such that $\tilde{\lambda} \in C(d)$, but $(\frac{1}{2} + \epsilon) \tilde{\lambda} \notin C(K, d)$ for any $\epsilon > 0$. In other words, *any given K*-partition may result in a 50% throughput loss for certain arrival density functions. Given a fixed *K*-partition, the idea behind this result is to find a set of points that can be scheduled simultaneously, but the related cells can not.

PROPOSITION 3.2. Let $K \in \mathbb{N}$, then there exists an arrival density function $\hat{\lambda}$ such that $\hat{\lambda} \in C(d)$, but $(\frac{1}{2} + \epsilon) \hat{\lambda} \notin C(K, d)$ for any $\epsilon > 0$.

PROOF. Assume that *Kd* is not an integer, and consider the set of points

$$\hat{\mathcal{P}} = \{ (kG, lG) : k, l = 1, 2, \dots, \lfloor 1/G \rfloor \}$$

with $G = (d + \lceil Kd \rceil / K)/2$. We further define an arrival density function

$$\hat{\lambda}(x,y) = \sum_{(\hat{x},\hat{y})\in\hat{\mathcal{P}}} \delta\left((x - \hat{x}, y - \hat{y}) \right).$$

It is readily seen that $\hat{\mathcal{P}}$ is a feasible set of points under the original reuse constraint d, so $\hat{\lambda} \in C(d)$. On the other hand, under K-partition, the point (kG, lG) belongs to cell $\mathcal{R}_{k\lceil Kd\rceil, l\lceil Kd\rceil}$. Thus the cell containing the point (kG, lG) interferes with the cell containing the point (k'G, l'G) if $|k - k'| + |l - l'| \leq 1$, which implies that $(\frac{1}{2} + \epsilon) \hat{\lambda} \notin C(K, d)$ for any $\epsilon > 0$.

To illustrate Proposition 3.2, consider the 9-partition shown in Figure 3.6 and assume the reuse distance is d = 0.35. It is easy to verify that the set of all points shown in the figure is a feasible subset according to the original reuse constraints, but the related cells are not interference-free. For example, cell $\mathcal{R}_{1,1}$ interferes with cell $\mathcal{R}_{5,1}$. Consequently, under region-based scheduling either all black points or all gray points can be scheduled at any time, but not both. Now assume flows of size $B \equiv 1$ uniformly arrive at the nine locations only at rate α (flows per location per time slot). A region-based scheduling strategy can only support any $\alpha \leq 1/2$, while any $\alpha < 1$ is within the network capacity region, by scheduling all locations simultaneously.

3.4 Concluding remarks

In this chapter we demonstrated that MaxWeight policies may fail to provide maximum stability in the presence of flow-level dynamics due to persistent spatial inefficiency. Loosely stated, MaxWeight policies tend to serve flows with large backlogs, even when the resulting spatial reuse is not particularly efficient, and fail to take advantage of maximum spatial reuse patterns involving flows with smaller backlogs.

While the root cause for instability observed in this chapter (flow-level dynamics) is the same as in Chapter 2, the way the presence of transient flows unhinges the MaxWeight scheduling algorithm is fundamentally different. Consequently, the remedies for instability discussed in [54, 55, 74] cannot be directly applied in the current setting, and the spatial inefficiency identified in this chapter calls for novel methods for stabilizing the system.

3.4 CONCLUDING REMARKS



Figure 3.6: The 9-partition with reuse distance d = 0.35.

We showed that the potential instability issues can be countered by traffic aggregation with sufficiently fine spatial granularity and adopting a region-based version of MaxWeight scheduling. A surprising fact is that the region-based approach involves a discretization with arrivals at a finite set of queues, which closely 'approximates' the arrivals in a continuum of locations as the spatial granularity increases, and yet the stability condition is markedly different. Even more remarkably, the set of admissible scheduling decisions is limited by the discretization, but the stability region for the MaxWeight strategy can be larger, i.e., constraining the set of feasible scheduling options can in fact expand the stability region of a specific scheduler. The complexity of region-based scheduling does not depend on the number of flows, in contrast to direct implementations of the algorithms in [54, 55, 74].

Finding the right granularity of the regions is non-trivial since the degree of traffic aggregation involves a trade-off between scheduling complexity, spatial efficiency, and network capacity. In particular, the suitable level of aggregation depends on the spatial load profile, and seems difficult to determine without explicit knowledge of the traffic parameters, thus detracting from one of the most appealing features of MaxWeight scheduling.

Proposition 3.2 establishes a negative result, in that no given region-based strategy can be expected to perform well for arbitrary spatial arrival densities. Note that the traffic pattern used in this counterexample is a discrete distribution which only injects traffic into the system at a finite number of locations. The question whether a universally stabilizing partitioning does exist when we restrict ourselves to a certain class of continuous arrival densities (e.g., smooth density functions) remains open, as well as the question whether the smoothness condition required for Theorem 3.1 can be removed.

SPATIAL INEFFICIENCY OF MAXWEIGHT SCHEDULING

Appendix

3.A Remaining proofs

3.A.1 Proof of Lemma 3.1

Let $\lambda \in C(d)$, $\mathcal{P} = \{X_1, X_2, \dots, X_n\} \in \Omega(d)$ and $K \ge 2\sqrt{2}/d$. We show that, with k_i, l_i such that $X_i \in \mathcal{R}_{k_i, l_i}, i = 1, \dots, n$,

$$\{(k_1, l_1), (k_2, l_2), \dots, (k_n, l_n)\} \in \Omega(K, d - 2\sqrt{2}/K).$$
(3.4)

That is, the points in \mathcal{P} belong to a feasible set of cells under *K*-partition with a reduced reuse distance $d - 2\sqrt{2}/K$. Consequently, any set of flows simultaneously scheduled under this strategy are located in a feasible set of cells under *K*-partition and reuse distance $d - 2\sqrt{2}/K$. Therefore this strategy is a legitimate region-based scheduling under *K*-partition with reuse distance $d - 2\sqrt{2}/K$, which means $\lambda \in C(K, d - 2\sqrt{2}/K)$.

We now prove (3.4). Let $i, j \in \{1, 2, ..., n\}$, $i \neq j$, and consider any two points $Y_i \in \mathcal{R}_{k_i, l_i}$ and $Y_j \in \mathcal{R}_{k_j, l_j}$. Then

$$\begin{aligned} \|Y_i - Y_j\| &= \|Y_i - X_i + X_i - X_j + X_j - Y_j\| \\ &\geq \|X_i - X_j\| - \|X_i - Y_i\| - \|X_j - Y_j\| \\ &\geq d - 2\sqrt{2}/K. \end{aligned}$$

As a result no two points $Y_i \in \mathcal{R}_{k_i, l_i}$ and $Y_j \in \mathcal{R}_{k_j, l_j}$ are within distance $d - 2\sqrt{2}/K$, $i \neq j$, and thus the subset of cells $\{(k_1, l_1), (k_2, l_2), \dots, (k_n, l_n)\}$ belongs to $\Omega(K, d - 2\sqrt{2}/K)$.

3.A.2 Proof of Lemma 3.2

Let $\omega \in \Omega(K, d - 2\sqrt{2}/K)$, and for $k, l = 1, 2, \dots, K/h$ denote,

$$\tilde{\mathcal{R}}_{k,l} = \left\{ (x, y) \in [0, 1]^2 : (x/h, y/h) \in \mathcal{R}_{k,l} \right\},\$$

the cells of size $(h/K)^2$ under (K/h)-partition. Consider two points $(x_1, y_1) \in \tilde{\mathcal{R}}_{k_1, l_1}$ and $(x_2, y_2) \in \tilde{\mathcal{R}}_{k_2, l_2}$, with $(k_1, l_1), (k_2, l_2) \in \omega$. It follows from the definition of $\Omega(K, d-2\sqrt{K})$ that $\|(x_1/h, y_1/h) - (x_2/h, y_2/h)\| \ge d-2\sqrt{2}/K$, which implies $\|(x_1, y_1) - (x_2, y_2)\| \ge h(d - 2\sqrt{2}/K) \ge d$, completing the proof.

3.A.3 Proof of Theorem 3.1

Let λ be a smooth arrival density function such that $(1 + \epsilon)\lambda \in C(d)$ for some $\epsilon > 0$. Lemma 3.1 then implies that for any $K > 2\sqrt{2}/d$, $(1 + \epsilon)\lambda \in C(K, d - 2\sqrt{2}/K)$, i.e., there exists $\pi(\omega) > 0$, $\omega \in \Omega(K, d - 2\sqrt{2}/K)$ with $\sum_{\omega \in \Omega(K, d - 2\sqrt{2}/K)} \pi(\omega) = 1$, such that

$$(1+\epsilon)\lambda_{k,l}\mathbb{E}[B] \le \sigma_{k,l} = \sum_{\omega \in \Omega(K, d-2\sqrt{2}/K)} \pi(\omega)\omega_{kl}$$
(3.5)

for all k, l = 1, 2, ..., K.

3.A REMAINING PROOFS

Now consider a randomized scheduling strategy which serves the set of cells in $\omega \in \Omega(K, d - 2\sqrt{2}/K)$ with probability $\pi(\omega)$. Let *h* be as in (3.3), and denote by $\tilde{\mathcal{R}}_{k,l}$ the cells under *K*/*h* partition. By Lemma 3.2 we know that any set $\omega \in$ $\Omega(K, d - 2\sqrt{2}/K)$ is valid under *K*/*h* partition and reuse distance *d*, and $\tilde{\mathcal{R}}_{k,l}$ is served a fraction of time $\sigma_{k,l}$.

Since the arrival density function λ is smooth, $\lambda(hx, hy)$ should be close to $\lambda(x, y)$ when *h* is close to 1. Specifically, it may be shown by the mean value theorem that

$$\lambda(hx, hy) \le \lambda(x, y) + 2\kappa^{(1)}(h-1),$$

The arrival intensity $\tilde{\lambda}_{k,l}$ of cell $\tilde{\mathcal{R}}_{k,l}$ can be bounded as

$$\begin{split} \tilde{\lambda}_{k,l} &= \int_{\tilde{\mathcal{R}}_{k,l}} \lambda(x, y) \, \mathrm{d}x \mathrm{d}y = h^2 \int_{\mathcal{R}_{k,l}} \lambda(hx, hy) \, \mathrm{d}x \mathrm{d}y. \\ &\leq h^2 \int_{\mathcal{R}_{k,l}} \left(\lambda(x, y) + 2\kappa^{(1)} (h-1) \right) \, \mathrm{d}x \mathrm{d}y. \end{split}$$
(3.6)

Now choose *K* large enough (and hence *h* small enough) such that $2\kappa^{(1)}(h-1) \le \epsilon \kappa^{(0)}/2$ and $h^2 \le \frac{1+\epsilon}{1+\epsilon/2}$, then

$$h^{2} \int_{\mathcal{R}_{k,l}} \left(\lambda(x, y) + 2\kappa^{(1)}(h-1) \right) dx dy$$

$$\leq (1+\epsilon) \int_{\mathcal{R}_{k,l}} \lambda(x, y) dx dy = (1+\epsilon)\lambda_{k,l}.$$
(3.7)

Combining (3.5)-(3.7) yields $\tilde{\lambda}_{kl} \mathbb{E}[B] \leq \sigma_{k,l}$ for all $k, l = 1, \dots, K/h$, i.e., $\lambda \in C(K/h, d)$.

50 SPATIAL INEFFICIENCY OF MAXWEIGHT SCHEDULING

4 Stability and insensitivity

In this chapter we divert our attention to the CSMA model, introduced in Section 1.3.2. This model is popular for its ability to provide accurate numerical [49] and qualitative [98] predictions, while retaining tractability; it will be used in Chapters 5-7 for a detailed study of the CSMA algorithm.

The CSMA model typically assumes that both the transmission durations and the back-off periods are exponentially distributed. In the first part of this chapter we show that the stationary distribution of the joint activity process in the CSMA network is insensitive with respect to the distribution of the back-off periods and the transmission durations. More precisely, the stationary distribution only depends on the mean back-off time and mean transmission duration. In the second part of this chapter we study the stability region of the unsaturated CSMA model discussed in Section 1.3.3, where packets are generated over time and buffers may occasionally empty. We consider a single-hop network in which packets immediately leave the network after transmission, and investigate the stability region of such networks. We will identify necessary and sufficient conditions for stability in the case of a complete conflict graph, and illustrate the difficulties that arise for general conflict graphs.

This chapter is organized as follows. In Section 4.1 we present a detailed model description and establish the insensitivity result for the stationary distribution of the model under saturated conditions. We then turn to an unsaturated scenario, and in Section 4.2 we present a necessary stability condition for general conflict graphs. In Section 4.3 the stability region is obtained for full conflict graphs, while Section 4.4 illustrates the difficulties that arise for partial conflict graphs. In Section 4.5 we make some concluding remarks.

STABILITY AND INSENSITIVITY

4.1 Insensitivity of the saturated model

In this section we consider the saturated CSMA model on a general conflict graph (V, E). We follow the description of this model in Section 1.3.2, with two exceptions: The transmission times of node *i* are independent and phase-type distributed with mean $1/\mu_i$, and the back-off periods of node *i* are independent and phase-type distributed with mean $1/\nu_i$. Recall that we denote $\sigma_i = \nu_i/\mu_i$.

We distinguish two scenarios, depending on whether or not the back-off period of a blocked node is frozen. In case the back-off period is frozen, it is resumed as soon as the node becomes unblocked again. When the back-off period of a node ends, it must be unblocked and will start a transmission. In case the back-off does not get frozen, the back-off period may end while the node is blocked, in which case the node simply starts a new back-off period. Under the distributional assumptions below, the probability of the back-off periods of two nodes ending simultaneously is equal to zero. Note that in case the back-off periods are exponentially distributed, it does not matter whether or not the back-off periods are frozen when a node becomes blocked, nor does it matter whether they are resumed or resampled when a node becomes unblocked again. However, for non-exponential back-off times and transmission durations these two scenarios are no longer equivalent.

As before we denote by $\Omega \subseteq \{0,1\}^n$ the set of all feasible joint activity states of the network. Let $X(t) \in \Omega$ represent the activity state of the network at time t, with $X_i(t)$ indicating whether node i is active $(X_i(t) = 1)$ at time t or not $(X_i(t) = 0)$. Since we consider generally distributed back-off times and transmission durations, $X_i(t)$ is no longer a Markov process. Denote by $\pi(x) = \lim_{t\to\infty} P\{X(t) = x\}$ the limiting probability that the joint activity state is $x \in \Omega$, assuming it exists.

We assume all back-off times and transmission durations to have phase-type distributions, and are interested in the stationary distribution of the Markov process that keeps track of the phase of each node. Using this stationary distribution, and the fact that phase-type distribution are dense in the space of all probability distributions with positive support (see [69, Chapter 2]), we will show that π is insensitive to the distributions of the back-off times and the transmission durations, and only depends on these distributions through the σ_i .

Let the back-off process of node *i* have a phase-type distribution with m_i +1 phases, where states 1,..., m_i are transient, and state m_i + 1 is absorbing. The corresponding starting probabilities are $\alpha_1, \ldots, \alpha_{m_i+1}$, and the transition rates are given by q_{kl} . Similarly, the transmission times of node *i* have a phase-type distribution with n_i + 1 phases (with state n_i + 1 absorbing), starting probabilities $\gamma_1, \ldots, \gamma_{n_i+1}$ and transition rates r_{kl} . Note that α_{m_i+1} (γ_{n_i+1}) represents the probability of a back-off period (transmission) of zero length.

Let β_k , $k = 1, ..., m_i$, represent the fraction of time that the back-off process of a node is in phase k, and let η_k , $k = 1, ..., n_i$, represent the fraction of time that the transmission process of a node is in phase k. The fractions β_k and η_k follow from Equations (4.8) and (4.9) in Appendix 4.A. We study the Markov process that keeps track of the activity of all nodes. Let $\omega = (\omega_1, ..., \omega_n)$ denote the state of the system. We use the convention that $\omega_i = -k$ when node i is in back-off phase k, and $\omega_i = k$ when node i is in the k-th transmission phase. So ω_i assumes values in $\Omega_i = \{-m_i, ..., -1\} \cup \{1, ..., n_i\}$, and the state space of the Markov process of the

4.1 INSENSITIVITY OF THE SATURATED MODEL

joint activity state is $\Omega_{\text{PH}} \subseteq \Omega_1 \times \cdots \times \Omega_n$. The stationary distribution π_{PH} of this Markov process has the following product-form solution:

LEMMA 4.1. Let the back-off times and transmission durations have a phase-type distribution. Then, regardless of back-off freezing,

$$\pi_{\rm PH}(\omega) = Z^{-1} \prod_{i:\omega_i \le -1} \beta_{-\omega_i} \prod_{i:\omega_i \ge 1} \sigma_i \eta_{\omega_i}, \quad \omega \in \Omega_{\rm PH}, \tag{4.1}$$

where Z is the normalization constant.

The proof of Lemma 4.1 is presented in Appendix 4.A.2. In this proof we treat the freezing and non-freezing systems in parallel, and we demonstrate that freezing indeed has no impact on the stationary distribution.

Using Lemma 4.1 we can now show that the node activity is insensitive to the distributions of the back-off times and transmission durations.

THEOREM 4.1. Let the back-off times and transmission durations have a phase-type distribution. Then, regardless of back-off freezing,

$$\pi(x) = Z^{-1} \prod_{i=1}^{n} \sigma_i^{x_i}, \quad x \in \Omega,$$

$$(4.2)$$

where Z is the normalization constant.

PROOF. Denote by $\Omega_{\text{PH}}(x)$ the set of all states $\omega \in \Omega_{\text{PH}}$ that correspond to $x \in \Omega$, i.e.,

$$\Omega_{\rm PH}(x) = \{ \omega \in \Omega_{\rm PH} \mid \forall i : \omega_i \le -1 \text{ if } x_i = 0, \ \omega_i \ge 1 \text{ if } x_i = 1 \}.$$

Then

$$\pi(x) = \sum_{\omega \in \Omega_{\text{PH}}(x)} \pi_{\text{PH}}(\omega) = Z^{-1} \sum_{\omega \in \Omega_{\text{PH}}(x)} \prod_{i:\omega_i \le -1} \beta_{\omega_i} \prod_{i:\omega_i \ge 1} \sigma_i \eta_{\omega_i} = Z^{-1} \prod_{i=1}^n \sigma_i^{x_i},$$

as $\sum_k \beta_k = 1$ and $\sum_k \eta_k = 1$.

In Theorem 4.1 we have shown that the limiting distribution of the activity process under generally distributed back-off times and transmission durations (4.2) is the same as the distribution for exponential distributions (1.8). This result was first shown in [97], and for generally distributed back-off periods and back-off freezing, partial proof arguments are presented in [49]. In the case without back-off freezing, the insensitivity result may be directly proven by representing the dynamics as those observed in an Engset network as considered in [8], Section 5.3. This model is essentially a variation of a Loss network (see Section 1.3.4) with a fixed and finite number of customers that are alternatively active and idle. So the number of idle customers trying to become active. The Engset network is constructed from the conflict graph (V, E): The links in the Engset network are of unit capacity and correspond to the undirected edges in the graph. Each node $i \in V$ is then represented as a customer in the Engset network which, when active, simultaneously uses links $\{i, j\}$ for all j

53

STABILITY AND INSENSITIVITY

such that $\{i, j\} \in E$. Each customer alternates between active and inactive phases. Finally, without back-off freezing, a customer who wishes to become active, is blocked if one of the required links is occupied (i.e., the corresponding node has an active neighbor), in which case the customer starts a new inactivity period (back-off). This corresponds to the so-called *jump-over* retrial behavior considered in [8], and insensitivity follows. In the case of back-off freezing, it does not seem possible to apply such a representation.

4.2 A necessary stability condition

In the previous section, we considered a scenario with saturated nodes that always have packets to transmit. We now turn to the unsaturated CSMA model from Section 1.3.3 where packets are generated over time and buffers may occasionally be empty. Recall that packets arrive at node *i* according to a renewal process with mean interarrival time $1/\lambda_i$. Nodes compete for access to the medium as before, with the modification that when unblocked nodes have no packets to transmit at the time a back-off period ends, they simply start a new back-off period. Once a packet has been transmitted, it leaves the system. The results in this section are valid irrespective of whether or not the back-off process is frozen.

Denote the traffic intensity at node *i* by $\rho_i = \lambda_i / \mu_i$, so that ρ_i is the fraction of time that this node has to be active in order to sustain the arrival rate λ_i . Define θ_i^* as the throughput of node *i*, i.e., the expected number of transmissions per unit of time, and denote the fraction of time that node *i* is active by τ_i^* , so that $\theta_i^* = \mu_i \tau_i^*$. Denote by $\tau_i = \sum_{x \in \Omega: x_i = 1} \pi(x)$ and $\theta_i = \mu_i \tau_i$ the fraction of time node *i* is active and throughput of node *i*, respectively, in the regime where all nodes are saturated. We have $\theta_i^* \leq \lambda_i$ by definition, with equality when node *i* is stable.

The next proposition provides a simple necessary condition for stability.

PROPOSITION 4.1. If $\lambda_i > \theta_i$ for all i = 1, ..., n, then all the nodes are unstable.

PROOF. We will show that all the nodes are unstable in the sense that each of the associated queues only empties finitely often, but it can in fact be established that the queue of node *i* grows in a linear manner at rate $\lambda_i - \theta_i$ in the long run. For convenience, we restrict the proof to a Poisson arrival process, but the arguments extend to general renewal arrival processes.

The idea behind the proof is as follows. The key observation is that once a queue empties, with non-zero probability the system may enter a state with all queues nonempty. Since all queues have positive drift in saturated conditions, all queues remain non-empty with non-zero probability. Thus, every time a queue empties, it may never do so again with probability bounded away from zero, and hence the queue only empties finitely often.

In order to formalize the above observation, let $Q_j(t)$ be the number of packets pending for transmission or in the process of being transmitted at node j at time t. Let $T_{i,n}$ be the time that the queue of node i empties for the n-th time. Let $U_{i,n} = \inf\{t \ge T_{i,n} : Q_j(t) \ge 1 \text{ for all } j = 1, ..., n\}$ be the first time after $T_{i,n}$ when all queues are non-empty. It is easily verified that there exists $b_1 > 0$ such that $\mathbb{P}(U_{i,n} < T_{i,n+1}) > b_1$ for all n. 4.2 A NECESSARY STABILITY CONDITION

Let $A_i(s, t)$ be the number of packet arrivals at node *j* during the time interval [s, t]and $B_i^*(s,t)$ the number of packet transmissions at node j during the time interval [s, t], so that

$$Q_j(t) = Q_j(s) + A_j(s,t) - B_j^*(s,t).$$
(4.3)

55

Moreover, denote by $B_i(U_{i,n}, t)$ the number of packet transmissions at node *j* during the time interval $[U_{i,n}, t]$ in a modified version of the network where the various nodes are in the exact same state at time $U_{i,n}$ and are all assumed to be saturated from that time onward. Define

$$V_{i,j,n} = \sup\{t : Q_j(s) \ge 1 \ \forall s \in [U_{i,n}, t]\},\$$

$$W_{i,j,n} = \sup\{t : A_j(U_{i,n}, s) - B_j(U_{i,n}, s) \ge 0 \ \forall s \in [U_{i,n}, t]\},\$$

and denote $V_{i,n} = \min_{j=1,\dots,n} V_{i,j,n}$ and $W_{i,n} = \min_{j=1,\dots,n} W_{i,j,n}$. By definition, $Q_j(t) \ge 1$ for all $t \in [U_{i,n}, V_{i,n}], j = 1, \dots, n$. Thus for all nodes j we have $B_i^*(U_{i,n}, t) = B_j(U_{i,n}, t)$ for all $t \in [U_{i,n}, V_{i,n}]$. From (4.3) with $s = U_{i,n}$ we see

$$Q_{j}(t) \geq 1 + A_{j}(U_{i,n}, t) - B_{j}^{*}(U_{i,n}, t) = 1 + A_{j}(U_{i,n}, t) - B_{j}(U_{i,n}, t), \quad t \in [U_{i,n}, V_{i,n}],$$

so $V_{i,n} \ge W_{i,n}$. Since $\frac{1}{t}\mathbb{E}[A_j(U_{i,n}, U_{i,n} + t)] = \lambda_j > \theta_j = \lim_{t \to \infty} \frac{1}{t}\mathbb{E}[B_j(U_{i,n}, U_{i,n} + t)]$, it follows that there exists $b_2 > 0$ such that $\mathbb{P}(V_{i,n} = \infty) \ge \mathbb{P}(W_{i,n} = \infty) > b_2$ for all n.

In conclusion, the probability that the queue of node *i* never empties again after it has emptied for the *n*-th time, is bounded from below by $b = b_1 b_2 > 0$. Thus, the total expected number of times that the queue of node *i* empties is bounded from above by $\sum_{n=0}^{\infty} (1-b)^n = 1/b$, which means that it only empties finitely often with probability 1. П

Proposition 4.1 establishes a connection between the throughput in the saturated model and stability in the unsaturated model. Recall that the saturation throughput follows directly from $\theta_i = \mu_i \sum_{x \in \Omega: x_i=1} \pi(x)$, with $\pi(\cdot)$ the limiting distribution in Theorem 4.1.

It might seem natural that a dual property to Proposition 4.1 holds as well, i.e., all the nodes are stable if $\lambda_i < \theta_i$ for all i = 1, ..., n. It is indeed the case that then at least one of the nodes must be stable, as otherwise the network behaves as in the saturated regime, and each node *i* would have a throughput $\theta_i^* = \theta_i$. This contradicts the fact that $\theta_i^* \leq \lambda_i < \theta_i$ for all nodes. However, it is not the case in general that *all* the nodes are stable if $\lambda_i < \theta_i$ for all i = 1, ..., n. In order to see that, we next consider an illustrative example.

4.2.1 Example: ring topology

Consider a 4-node ring topology, i.e., n = 4 and $E = \{\{1, 2\}, \{2, 3\}, \{3, 4\}, \{4, 1\}\}$. Suppose that $v_i \equiv v$ and $\mu_i \equiv 1$, so that $\theta_i \equiv \theta = v(1+v)/(1+4v+2v^2)$. Also, let $\lambda_1 = \lambda_2 = \lambda_3 \equiv \lambda < \theta$, and $\lambda_4 = 0$. Assume that the arrival processes are Poisson and that the back-off periods and transmission durations are exponentially distributed.

First observe that both nodes 1 and 3 must be stable. In case either of these nodes were unstable, the fraction of time it would be active is bounded from below by

$$(1 - \tau_2^*)\frac{\nu}{1 + \nu} \ge (1 - \lambda)\frac{\nu}{1 + \nu} \ge \left(1 - \frac{\nu(1 + \nu)}{1 + 4\nu + 2\nu^2}\right)\frac{\nu}{1 + \nu} > \frac{\nu(1 + \nu)}{1 + 4\nu + 2\nu^2} = \theta > \lambda,$$

STABILITY AND INSENSITIVITY

which yields a contradiction.

A 'busy period' of node 2 is said to begin when node 2 starts a transmission after at least one transmission of nodes 1 and/or 3 and to end when node 2 completes a transmission which is followed by at least one transmission of nodes 1 and/or 3. The time period *U* between two successive busy periods of node 2 is referred to as an 'idle period' of node 2. Denote by *V* the amount of time from the start of a busy period of node 2 until the first packet arrival at node 1 or 3 and by *W* the possible remaining transmission time of node 2 after *V*. The number of further transmissions of node 2 after *W* during that busy period is bounded from above by a geometrically distributed random variable with parameter 1/2. Thus the total amount of time *T* that node 2 is active during the busy period satisfies $\mathbb{E}[T] \leq \mathbb{E}[V] + \mathbb{E}[W] + 1 \leq \frac{1}{2\lambda} + 2$. Now distinguish two cases: (i) $\mathbb{E}[U] \geq \mathbb{E}[T] + 1$; and (ii) $\mathbb{E}[U] < \mathbb{E}[T] + 1$, and denote by *S* the amount of time that node 2 is idle during the busy period. In case (i), we have

$$\tau_2^* = \frac{\mathbb{E}[T] - \mathbb{E}[S]}{\mathbb{E}[T] + \mathbb{E}[U] + \mathbb{E}[S]} \le \frac{\mathbb{E}[T]}{\mathbb{E}[T] + \mathbb{E}[U]} \le \frac{\frac{1}{2\lambda} + 2}{\frac{1}{\lambda} + 5} = \frac{1 + 4\lambda}{2 + 10\lambda}$$

It is easily verified that the latter upper bound is less than λ when $\lambda > (1 + \sqrt{11})/10$, and thus node 2 must be unstable when $\lambda > (1 + \sqrt{11})/10$.

In case (ii), note that by the time the last transmission of either node 1 or node 3 during an idle period of node 2 ends, one of these two has been inactive for at least an expected amount of time 1. Thus the fraction of time v that either nodes 1 and 2 or nodes 2 and 3 are inactive, is bounded from below as

$$\upsilon \ge \frac{1 + \mathbb{E}[S]}{\mathbb{E}[T] + \mathbb{E}[U] + \mathbb{E}[S]} \ge \frac{1}{\mathbb{E}[T] + \mathbb{E}[U]} \ge \frac{1}{\frac{1}{\lambda} + 5} = \frac{\lambda}{1 + 5\lambda}$$

where the second inequality follows from the fact that $\mathbb{E}[T] + \mathbb{E}[U] \ge 1$. Denote by v_{12} and v_{23} the fraction of time that nodes 1 and 2 are inactive, and the fraction of time that nodes 2 and 3 are inactive, respectively. Then

$$\min\{\tau_1^* + \tau_2^*, \tau_2^* + \tau_3^*\} = \min\{1 - \upsilon_{12}, 1 - \upsilon_{23}\} = 1 - \max\{\upsilon_{12}, \upsilon_{23}\}$$
$$\leq 1 - \frac{1}{2}(\upsilon_{12} + \upsilon_{23}) \leq 1 - \frac{1}{2}\upsilon \leq \frac{2 + 9\lambda}{2 + 10\lambda}.$$

It is easily verified that the latter upper bound is less than 2λ when $\lambda > (5 + \sqrt{185})/40$. Since nodes 1 and 3 are stable, it follows that node 2 must be unstable when $\lambda > (5 + \sqrt{185})/40$.

We conclude that in either case node 2 is unstable when $\lambda > (5 + \sqrt{185})/40 < 0.47$, regardless of the value of ν . However, θ approaches 1/2 for ν sufficiently large, and hence the condition $\lambda < \theta$ is not sufficient for stability.

Figure 4.1 shows the stability region (obtained by simulation) of the 4-node ring network as a function of λ_2 and λ_4 , for $\mu = 1$ and $\nu = 10$. For arrival rates outside of the area demarcated by the solid line, at least one node is unstable. The area enclosed by the dashed line represents all rate vectors such that max{ λ_2, λ_4 } < $\theta = 110/241 \approx 0.456$. We compare $\lambda_1 = \lambda_3 = 0.3$ (Figure 4.1(a)) and $\lambda_1 = \lambda_3 = 0.45$ (Figure 4.1(b)), so in both cases we have max{ λ_1, λ_3 } < θ . In Figure 4.1(a) we see that max{ λ_2, λ_4 } < θ is sufficient for stability. However, from Figure 4.1(b) it is clear that as λ_1 and λ_3

4.3 STABILITY FOR FULL CONFLICT GRAPHS

increase, the stability region decreases in size and no longer includes all rate vectors such that $\max{\lambda_2, \lambda_4} < \theta$.

Observe that the range of values of λ_2 that stabilize the system grows with λ_4 , so nodes 2 and 4 benefit from each other's activity. This was already hinted at in the above instability example, where it was shown that node 2 may become unstable when node 4 is removed from the network. Indeed, increasing activity from nodes 2 and 4 forces nodes 1 and 3 to operate in a more efficient fashion, i.e., simultaneous activity, thus increasing spatial reuse.



Figure 4.1: The stability region of a 4-node ring network.

4.3 Stability for full conflict graphs

In Section 4.2 we have seen that although necessary stability conditions can be obtained by considering the saturated model, the case of a ring network already gives rise to intricate stability conditions. The case of a full conflict graph is considerably simpler, and the stability conditions can be explicitly derived, as is shown next. In the remainder, we restrict ourselves to the case of back-off freezing.

We assume that all nodes mutually interfere, so at most one node can be active at any time. Without loss of generality assume that the nodes are ordered such that

$$\frac{\lambda_1}{\nu_1} \leq \frac{\lambda_2}{\nu_2} \leq \cdots \leq \frac{\lambda_n}{\nu_n}.$$

Recall that $\rho_i = \lambda_i / \mu_i$, and denote

$$\hat{\tau}_i = \frac{\sigma_i}{1 + \sum_{j=i}^n \sigma_j} \left(1 - \sum_{j=1}^{i-1} \rho_j \right).$$

These $\hat{\tau}_i$ may be interpreted as the fraction of time that node *i* is active, assuming that nodes $1, \ldots, i-1$ are all stable, while nodes i, \ldots, n are all saturated. Also, define

STABILITY AND INSENSITIVITY

 $i_{\max} = \max\{i \in \{1, ..., n\} : \rho_i < \hat{\tau}_i\}$, with the convention that $i_{\max} = 0$ when $\rho_i > \hat{\tau}_i$ for all i = 1, ..., n, and assume $\rho_{i_{\max}+1} > \hat{\tau}_{i_{\max}+1}$ in case $i_{\max} < n$. The interpretation of $\hat{\tau}_i$ suggests that node i_{\max} is the stable node with the highest index, as will be shown in the following theorem.

THEOREM 4.2. Nodes $1, \ldots, i_{max}$ are stable, while nodes $i_{max} + 1, \ldots, n$ are unstable.

PROOF. For compactness, denote by τ_0^* the fraction of time that all nodes are inactive. As noted in Section 4.2, we have that $\theta_i^* \leq \lambda_i$, with equality when node *i* is stable, and thus $\tau_i^* \leq \rho_i$, with equality when node *i* is stable. In view of the back-off freezing, the back-off process of node *i* is only running when all nodes are inactive, and hence we have $\tau_i^* = \gamma_i/\nu_i$, with γ_i the expected number of back-offs of node *i* per unit of time. (Without back-off freezing, this relationship still holds for exponential back-off time distributions, but for general back-off time distributions there does not seem to be a simple connection between γ_i and τ_0^* in that case.) By definition, the probability that node *i* has a packet to transmit when a back-off period ends, equals $p_i = \theta_i^*/\gamma_i$. Combining these two relationships, we obtain the identity $\theta_i^* = p_i \nu_i \tau_0^*$ and thus $\tau_i^* = p_i \sigma_i \tau_0^*$. In particular, $\tau_i^* = \sigma_i \tau_0^*$ when node *i* is unstable. Hence $\rho_i \leq \sigma_i \tau_0^*$, i.e., $\lambda_i \leq \nu_i \tau_0^*$, when node *i* is stable, while $\rho_i \geq \sigma_i \tau_0^*$, i.e., $\lambda_i \geq \nu_i \tau_0^*$ when node *i* is stable nodes is of the form $\{1, \ldots, i^*\}$ for some $i^* \in \{0, \ldots, n\}$. It remains to be shown that $i^* = i_{max}$.

First observe that $\sum_{i=0}^{n} \tau_i^* = 1$, $\tau_i^* = \rho_i$ for all $i = 1, ..., i^*$, and $\tau_i^* = \sigma_i \tau_0^*$ for all $i = i^* + 1, ..., n$. This yields

$$\tau_0^* = \frac{1}{1 + \sum_{j=i^*+1}^n \sigma_j} \Big(1 - \sum_{i=1}^{i^*} \rho_i \Big).$$

Further observe the equivalence relation

$$\rho_i > \frac{\sigma_i}{1 + \sum_{j=i}^n \sigma_j} \left(1 - \sum_{j=1}^{i-1} \rho_j \right) \iff \rho_i + \rho_i \sum_{j=i}^n \sigma_j > \sigma_i - \sigma_i \sum_{j=1}^{i-1} \rho_j$$
$$\implies \rho_i + \rho_i \sum_{j=i+1}^n \sigma_j > \sigma_i - \sigma_i \sum_{j=1}^i \rho_j \iff \rho_i > \frac{\sigma_i}{1 + \sum_{j=i+1}^n \sigma_j} \left(1 - \sum_{j=1}^i \rho_j \right).$$

Since the nodes are indexed such that $\rho_i/\sigma_i = \lambda_i/\nu_i \le \rho_{i+1}/\sigma_{i+1} = \lambda_{i+1}/\nu_{i+1}$, we obtain the property

$$\rho_i > \hat{\tau}_i^* \Longrightarrow \rho_{i+1} > \hat{\tau}_{i+1}. \tag{4.4}$$

Now suppose that $0 \le i^* < i_{\max}$. The fact that node $i^* + 1 \le n$ is unstable means that

$$\rho_{i^*+1} > \tau_{i^*+1}^* = \sigma_{i^*+1}\tau_0^* = \frac{\sigma_{i^*+1}}{1 + \sum_{j=i^*+1}^n \sigma_j} \left(1 - \sum_{i=1}^{i^*} \rho_i\right) = \hat{\tau}_{i^*+1}.$$

Property (4.4) then implies that $\rho_i > \hat{\tau}_i$ for all $i = i^* + 1, ..., n$, which contradicts $i_{\max} \ge i^* + 1$, and hence we must have $i^* \ge i_{\max}$.

The fact that node i^* is stable means that

$$\rho_{i^*} = \tau_{i^*}^* \le \sigma_{i^*} \tau_0^* = \frac{\sigma_{i^*}}{1 + \sum_{j=i^*+1}^n \sigma_j} \Big(1 - \sum_{i=1}^{i^*} \rho_i \Big).$$

4.3 STABILITY FOR FULL CONFLICT GRAPHS

Property (4.4) then implies that

$$\rho_{i^*} \leq \frac{\sigma_{i^*}}{1 + \sum_{j=i^*}^n \sigma_j} \left(1 - \sum_{i=1}^{i^*-1} \rho_i \right) = \hat{\tau}_{i^*},$$

and hence we must have $i_{\text{max}} \ge i^*$.

The result in Theorem 4.2 in fact holds for any stationary traffic process as long as the service is infinitely divisible, with the ρ_i values representing the mean amount of traffic generated per time unit as measured in units of transmission time. The form of the stability conditions is rather reminiscent of those for polling systems with k_i -limited or Weighted Fair Queueing (WFQ) service disciplines [25] and Generalized Processor Sharing (GPS) queues [13, 46].

Noting that $i_{\text{max}} = n$ if and only if $\rho_i < \hat{\tau}_i$ for all i = 1, ..., n, Theorem 4.2 in particular gives the following necessary and sufficient condition for all nodes to be stable.

COROLLARY 4.1. All nodes are stable if and only if $\rho_i < \hat{\tau}_i$ for all i = 1, ..., n.

The explicit and relatively simple form of the stability condition established in Corollary 4.1 is highly remarkable as it starkly contrasts with those for slotted Aloha systems, which even for a complete conflict graph with three or more nodes have remained largely elusive, see for instance [3, 85, 86] for bounds and partial results.

The next result shows that for full conflict graphs, the dual property of Proposition 4.1 does hold (which in general is not the case; see Section 4.2.1).

COROLLARY 4.2. All nodes are stable if $\lambda_i < \theta_i$ for all i = 1, ..., n.

PROOF. For full conflict graphs we have that $\tau_i = \sigma_i / (1 + \sum_{i=1}^n \sigma_i)$. Note that $\lambda_i < \theta_i$ for all i = 1, ..., n implies $\rho_n < \tau_n$ and $\tau_n = (1 - \tau_1 - \cdots - \tau_{n-1}) \frac{\sigma_n}{1 + \sigma_n} < (1 - \rho_1 - \cdots - \rho_{n-1}) \frac{\sigma_n}{1 + \sigma_n} = \hat{\tau}_n$. This yields $i_{\text{max}} = n$, which completes the proof.

We conclude this section by two further consequences of Theorem 4.2 that are helpful when only the total load $\sum_{i=1}^{n} \rho_i$ is known. Denote $\sigma_{\min} = \min_{i=1,...,n} \sigma_i$ and $\sigma_{\max} = \max_{i=1,...,n} \sigma_i$.

COROLLARY 4.3. All nodes are stable if $\sum_{i=1}^{n} \rho_i < \frac{\sigma_{\min}}{1+\sigma_{\min}}$. This condition is sharp in the sense that if $\rho_1 = \cdots = \rho_{n-1} = 0$ and $\sigma_n = \sigma_{\min}$, then $\rho_n < \frac{\sigma_{\min}}{1+\sigma_{\min}}$ is necessary for node n to be stable.

COROLLARY 4.4. At least one node is unstable if $\sum_{i=1}^{n} \rho_i > \frac{\sigma_{\max}}{1/n + \sigma_{\max}}$. This condition is sharp in the sense that if $\rho_1 = \cdots = \rho_n = \rho$ and $\sigma_1 = \cdots = \sigma_n = \sigma$, then $\rho < \frac{\sigma}{1+n\sigma}$ is sufficient for all nodes to be stable.



STABILITY AND INSENSITIVITY

4.4 Stability for partial conflict graphs

In Section 4.3 we focused on the case of a full conflict graph, and derived explicit necessary and sufficient conditions for stability. In this section we allow for partial conflict graphs, and will show that in general the stability conditions cannot be represented in such an explicit form. In particular, we illustrate the difficulties that arise in star topologies, and then argue that all (non-complete) graphs contain such a star topology as an induced subgraph.

4.4.1 Star topologies

Consider a star topology, where the leaf nodes 1, ..., n - 1 all interfere with the root node n, but not with each other, i.e., $E = \{1, ..., n - 1\} \times \{n\}$. The stability region may then be characterized by:

$$\rho_i < (1 - \rho_n) \frac{\sigma_i}{1 + \sigma_i}, \qquad i = 1, \dots, n - 1,$$
(4.5)

and

$$\rho_n < \hat{\tau}_n, \tag{4.6}$$

with $\hat{\tau}_n$ representing the fraction of time that node *n* would be active if it were saturated.

By definition, inequality (4.6) is necessary and sufficient for the root node n to be stable, and given that node n is stable, the inequalities (4.5) are necessary and sufficient for all the leaf nodes 1, ..., n - 1 to be stable as well. The boundary of the stability region consists of a total of n segments, with n - 1 linear segments defined by the inequalities (4.5), where the corresponding leaf node is critically loaded, and 1 segment which is not likely to be linear in general, described by the inequality (4.6), where the root node is critically loaded.

There does not seem to be a closed-form expression available for $\hat{\tau}_n$ in general, in fact not even for n = 3, so that the inequality (4.6) is not so explicit. The next lemma however provides a useful closed-form lower bound for $\hat{\tau}_n$.

COROLLARY 4.5. Assuming exponential back-off times and transmission durations, we have $\hat{\tau}_n \geq \tau_n$, with

$$\tau_n = \frac{\sigma_n}{\sigma_n + \prod_{i=1}^{n-1} (1 + \sigma_i)}.$$

PROOF. Noting that the star topology is a bipartite graph, the statement of the corollary follows directly from Proposition 4.2 (presented in Appendix 4.A), with $V_1 = \{n\}$ and $V_2 = \{1, ..., n - 1\}$.

Corollary 4.5 implies that as long as $\rho_n < \tau_n$, the root node *n* is guaranteed to be stable, and thus the conditions (4.5) are necessary and sufficient for all the leaf nodes $1, \ldots, n-1$ to be stable as well. Noting that $\tau_i = (1 - \tau_n) \frac{\sigma_i}{1 + \sigma_i}$, the inequalities (4.5) may be expressed as

$$\frac{\rho_i}{\tau_i} < \frac{1-\rho_n}{1-\tau_n}, \qquad i=1,\ldots,n-1,$$

4.4 STABILITY FOR PARTIAL CONFLICT GRAPHS

or

$$\max_{i=1,\dots,n-1}\frac{\rho_i}{\tau_i}<\frac{1-\rho_n}{1-\tau_n},$$

or

$$\max_{i=1,\dots,n-1}\frac{\rho_i(1+\sigma_i)}{\sigma_i} < 1-\rho_n.$$

Thus simple necessary and sufficient conditions for stability arise when the traffic load of the root node *n* is sufficiently low, in particular when the ratio of the traffic load ρ_n to the saturated throughput τ_n is relatively low compared to that of the leaf nodes 1, ..., n - 1. Specifically, suppose that

$$\frac{\rho_n}{\tau_n} \leq \max_{i=1,\dots,n-1} \frac{\rho_i}{\tau_i}$$

and let $j_{\max} = \arg \max_{j=1,\dots,n-1} \frac{\rho_j}{\tau_j}$. Then

$$\frac{\rho_{j_{\max}}}{\tau_{j_{\max}}} < 1 - \rho_n$$

is a simple necessary and sufficient condition for stability of all nodes. In order to prove that, it suffices to show that the latter condition implies that $\rho_n < \tau_n$. Suppose that is not the case, then $\max_{i=1,...,n-1} \frac{\rho_i}{\tau_i} \ge 1$, so

$$\max_{i=1,\dots,n-1} \frac{\rho_i}{\tau_i} < \frac{1-\rho_n}{1-\tau_n} \Rightarrow \frac{1-\rho_n}{1-\tau_n} > 1.$$

In particular, if $\frac{\rho_n}{\tau_n} = \max_{i=1,...,n-1} \frac{\rho_i}{\tau_i}$, then $\rho_n < \tau_n$ is a simple necessary and sufficient condition for the stability of all nodes.

A further observation is that $\rho_i < \tau_i$, i = 1, ..., n, is a sufficient condition for all nodes 1, ..., n to be stable. It might seem that this is a trivial fact, which in fact should hold for any conflict graph, but that is not the case as was illustrated in the counterexample in Section 4.2.1. However, an explicit formulation of the *necessary* and sufficient stability condition for star networks stays beyond our reach, as $\hat{\tau}_n$ in (4.6) remains elusive.

4.4.2 Stability conditions for general graphs

The fact that an explicit condition for stability in the star network appears elusive, illustrates the difficulty of obtaining the stability condition for general networks. Indeed, an explicit characterization of the stability region is difficult for any network that is not a complete graph, for which the stability condition was explicitly obtained in Theorem 4.2.

One way to argue this is to show that any network is either a complete graph or contains a 3-node star network as a subgraph. In order to see that, we may focus on a connected graph. Consider an arbitrary node in the graph, say node 1, as well as the set C_1 of all of its neighbors. If one of the nodes in the set C_1 has a neighbor that is not node 1 or in the set C_1 , then this induces a 3-node star network. Otherwise, node 1 along with the nodes in the set C_1 make up the entire graph (since the graph is connected). If the nodes in the set C_1 are not fully connected, then this induces a

STABILITY AND INSENSITIVITY



Figure 4.2: A complete graph and a 3-node star subgraph.

3-node star network again. Otherwise, node 1 along with the nodes in the set C_1 are fully connected, and the graph is complete. This argument is illustrated in Figure 4.2.

Hence, the fact that the 3-node star network (for which the stability condition does not seem to admit an explicit characterization) is contained in every non-complete network, provides strong indication of the hardness of the problem of determining the stability region for general networks. That is, characterizing the set of traffic vectors (ρ_1, \ldots, ρ_n) for which the system is stable is challenging for given back-off rates ν_1, \ldots, ν_n .

In contrast, determining whether there exist back-off rates v_1, \ldots, v_n for which the system is stable for a given traffic vector (ρ_1, \ldots, ρ_n) is relatively easy. Specifically, in case of a star topology, such back-off rates exist if and only if $\rho_n + \max_{i=1,\ldots,n-1} \rho_i < 1$. In general, such back-off rates exist if and only if the traffic vector $(\rho_1, \ldots, \rho_n) < \hat{\rho}$ for some $\hat{\rho} \in \operatorname{conv}(\Omega)$, see [32]. The latter property in fact serves as the basis for the adaptive CSMA algorithms discussed in Section 1.3.5.

4.5 Concluding remarks

In this chapter we examined the insensitivity and stability of CSMA networks. We proved that when all nodes are saturated, the limiting behavior of these networks is insensitive to the distribution of both the transmission durations and the back-off times. More precisely, the stationary distribution of the activity process only depends on the distributions of the back-off times and transmission durations through their means. The insensitivity holds irrespective of whether or not an active node freezes the back-off process of neighboring nodes. The CSMA model considered in Chapters 5-7 assumes exponential distributions, but since the analysis in these chapters take the stationary distribution of the activity process as a starting point, insensitivity implies that all results obtained in Chapters 5-7 hold for generally distributed back-off times and transmission durations.

We then turned to a situation where nodes are subject to packet dynamics, and established a simple necessary condition for stability for general conflict graphs. Explicit necessary and sufficient stability conditions were derived for the case of complete conflict graphs. Moreover, we illustrated the difficulty of deriving similar conditions for partial conflict graphs.

4.5 CONCLUDING REMARKS

In the above analysis, we used a continuous-time model to capture CSMA dynamics. Alternatively we may consider a model where time is slotted and neighboring nodes occasionally start transmitting in the same slot, leading to a collision. Extending our results to the case where collisions may occur is difficult. If we assume that when two nodes involved in a collision start and end the corresponding transmissions at the same epochs, then insensitivity can be shown using the Engset network representation (see the discussion following Theorem 4.1), but only in the case without back-off freezing. It remains unclear whether insensitivity holds when time is slotted in the case with back-off freezing. Regarding stability issues, adding collisions greatly complicates the analysis even in the case of networks with full interference (similar to the classical problem of deriving network stability conditions under Aloha protocols). In Section 4.3, the absence of collisions simplifies the stability analysis, because in this case, packets are successfully transmitted whenever the medium is busy.
STABILITY AND INSENSITIVITY

Appendix

4.A Auxiliary results and remaining proofs

4.A.1 A stochastic comparison result

Consider a bipartite graph such that $V = V_1 \cup V_2$, with $V_1 \cap V_2 = \emptyset$ and $E \subseteq V_1 \times V_2$. We will show that in the situation where all the nodes are saturated, the throughputs of the nodes in V_1 and V_2 are lower and higher respectively, than in case only the nodes in V_1 are saturated. Let $V_1 \subseteq \{1, ..., n\}$ and define $\theta_i^*(V_1)$ as the throughput of node *i* in the situation where the nodes in V_1 are saturated.

PROPOSITION 4.2. Assuming exponential back-off times and transmission durations, we have $\theta_i^*(V_1) \ge \theta_i$ for all $i \in V_1$ and $\theta_i^*(V_1) \le \theta_i$ for all $i \in V_2$.

PROOF. The proof relies on stochastic coupling [64]. Let $N_i^{\lambda}(t)$, $i \in V_2$, $N_i^{\mu}(t)$, i = 1, ..., n, and $N_i^{\nu}(t)$, i = 1, ..., n, be independent Poisson processes of rates λ_i , μ_i , and ν_i , respectively. We will use these Poisson processes to construct processes $X^*(t) = (X_1^*(t), ..., X_n^*(t))$ and $Y(t) = (Y_i(t))_{i \in V_2}$, representing the activity process and the queue length process in the scenario where the nodes $i \in V_1$ are saturated, and $X(t) = (X_1(t), ..., X_n(t))$ representing the activity process in case all the nodes are saturated. It is easily verified that viewed in isolation, the processes $X^*(t)$ and Y(t) as constructed above obey the same statistical laws as the activity process and the queue length process in the scenario where the nodes $i \in V_1$ are saturated, while the process X(t) is governed by the same statistical laws as the activity process in case all the nodes are saturated.

We assume that $X_i^*(0) \ge X_i(0)$ for all $i \in V_1$ and $X_i^*(0) \le X_i(0)$ for all $i \in V_2$, and allow $Y_i(0)$, $i \in V_2$, to be arbitrary. We will prove that $X_i^*(t) \ge X_i(t)$ for all $i \in V_1$ and $X_i^*(t) \le X_i(t)$ for all $i \in V_2$. Since the stationary distribution of the processes $X^*(t)$ and X(t) does not depend on the initial state, and $\theta_i^*(V_1) = \mu_i \mathbb{E}[X_i^*] = \mu_i \mathbb{P}(X_i^* = 1)$ and $\theta_i = \mu_i \mathbb{E}[X_i] = \mu_i \mathbb{P}(X_i = 1)$, the statement of the proposition then follows.

We prove the above inequalities by induction. Let t be a time epoch at which an event occurs in one of the Poisson processes. We will show that if the inequalities hold at time t^- , that they then continue to hold at time t^+ . We distinguish three cases, depending on in which of the various Poisson processes the event occurs.

We first consider an event in the process $N_i^{\lambda}(t)$, reflecting a packet arrival at one of the unsaturated nodes $i \in V_2$. In that case, we set $Y_i(t^+) = Y_i(t^-) + 1$. Note that the values of $X_i^*(t)$ and $X_i(t)$ are not affected, and hence the inequalities trivially continue to be valid. Second, we consider an event in the process $N_i^{\mu}(t)$, corresponding to a potential transmission completion at node *i*. In that case, we set $X_i^*(t^+) = X_i(t^+) = 0$, and in case $i \in V_2$,

$$Y_i(t^+) = Y_i(t^-) - X_i^*(t^-), (4.7)$$

reflecting a potential packet departure. Since $X_i^*(t^+) = X_i(t^+) = 0$, the inequalities remain trivially satisfied. Third, we consider an event in the process $N_i^{\nu}(t)$, corresponding to a potential activation at node *i*. In that case we set $X_i(t^+) = 1$ if $X_j(t^-) = 0$ for all $j \in C_i$, with C_i representing the set of neighbors of node *i*. Moreover, in case $i \in V_1$, we set $X_i^*(t^+) = 1$ if $X_j^*(t^-) = 0$ for all $j \in C_i$, while in case $i \in V_2$, we set $X_i^*(t^+) = 1$ if $Y_i(t^-) \ge 1$ and $X_j^*(t^-) = 0$ for all $j \in C_i$.

4.A AUXILIARY RESULTS AND REMAINING PROOFS

The fact that for $i \in V_1$, $C_i \subseteq V_2$, and $X_j^*(t^-) \leq X_j(t^-)$ for all $j \in V_2$ implies that $X_i(t^+) = 1$ forces $X_i^*(t^+) = 1$. Likewise, the fact that for $i \in V_2$, $C_i \subseteq V_1$, and $X_j^*(t^-) \geq X_j(t^-)$ for all $j \in V_1$ implies that $X_i^*(t^+) = 1$ forces $X_i(t^+) = 1$. Hence, the inequalities continue to hold. Also, note that $X_i^*(t) = 0$ whenever $Y_i(t) = 0$, or equivalently, $X_i^*(t) = 1$ can only occur when $Y_i(t) \geq 1$, so that (4.7) leaves $Y_i(t) \geq 0$ for all $t \geq 0$.

4.A.2 Proof of Lemma 4.1

We have that

$$\beta_{u}\left(q_{u,m_{i}+1}\left(1-\frac{\alpha_{u}}{1-\alpha_{m_{i}+1}}\right)+\sum_{\substack{l=1\\l\neq u}}^{m_{i}}q_{u,l}\right)=\sum_{\substack{l=1\\l\neq u}}^{m_{i}}\beta_{l}\left(q_{l,u}+q_{l,m_{i}+1}\frac{\alpha_{u}}{1-\alpha_{m_{i}+1}}\right),$$
(4.8)

with $\sum_{u=1}^{m_i} \beta_u = 1$, and

$$\eta_u \left(r_{u,n_i+1} \left(1 - \frac{\gamma_u}{1 - \gamma_{n_i+1}} \right) + \sum_{\substack{l=1\\l \neq u}}^{n_i} r_{u,l} \right) = \sum_{\substack{l=1\\l \neq u}}^{n_i} \eta_l \left(r_{l,u} + r_{l,n_i+1} \frac{\gamma_u}{1 - \gamma_{n_i+1}} \right), \quad (4.9)$$

with $\sum_{u=1}^{n_i} \eta_u = 1$. It is readily seen that

$$\nu_i = (1 - \alpha_{m_{i+1}})^{-1} \sum_{u=1}^{m_i} \beta_u q_{u,m_i+1}, \qquad (4.10)$$

$$\mu_i = (1 - \gamma_{n_{i+1}})^{-1} \sum_{u=1}^{n_i} \eta_u r_{u,n_i+1}.$$
(4.11)

In order to show that the π_{PH} in (4.1) is indeed the limiting distribution of the Markov process of interest, it suffices to show that π_{PH} satisfies the global balance equations of this process. However, rather than doing this directly, we study for each node *i* the partial balance equations that equate the rate into and out of a state by changes to this node only. As the global balance equations can be obtained by summing these partial balance equations over all nodes, it is sufficient to show that π_{PH} satisfies all partial balance equations.

Let C_i denote the set of neighbors of node *i*, and define $C_i^+ = C_i \cup \{i\}$. Let $T_k^i(\omega)$ denote the operator that changes the *i*-th component of ω to *k*, while leaving the other components intact. When node *i* is inactive and unblocked we see the following transitions to node *i* (irrespective of freezing) $\forall \omega$ s.t. $\omega_j \leq -1 \ \forall j \in C_i^+$

$$\pi_{\rm PH}(\omega) \left(q_{\omega_{i},m_{i}+1} \left(1 - \frac{\alpha_{\omega_{i}} \gamma_{n_{i}+1}}{1 - \alpha_{m_{i}+1} \gamma_{n_{i}+1}} \right) + \sum_{\substack{k=1\\k\neq\omega_{i}}}^{m_{i}} q_{\omega_{i},k} \right)$$

$$= \sum_{k=1}^{n_{i}} \pi_{\rm PH}(T_{k}^{i}(\omega)) r_{k,n_{i}+1} \frac{\alpha_{\omega_{i}}}{1 - \alpha_{m_{i}+1} \gamma_{n_{i}+1}}$$

$$+ \sum_{\substack{k=1\\k\neq\omega_{i}}}^{m_{i}} \pi_{\rm PH}(T_{-k}^{i}(\omega)) \left(q_{k,\omega_{i}} + q_{k,m_{i}+1} \frac{\alpha_{\omega_{i}} \gamma_{n_{i}+1}}{1 - \alpha_{m_{i}+1} \gamma_{n_{i}+1}} \right).$$
(4.12)

STABILITY AND INSENSITIVITY

If node *i* is active then we have (irrespective of freezing) $\forall \omega$ s.t. $\omega_i \ge 1$

$$\pi_{\rm PH}(\omega) \left(r_{\omega_{i},n_{i}+1} \left(1 - \frac{\alpha_{m_{i}+1} y_{\omega_{i}}}{1 - \alpha_{m_{i}+1} y_{n_{i}+1}} \right) + \sum_{\substack{k=1\\k\neq\omega_{i}}}^{n_{i}} r_{\omega_{i},k} \right)$$

$$= \sum_{k=1}^{m_{i}} \pi_{\rm PH}(T_{-k}^{i}(\omega)) q_{k,m_{i}+1} \frac{y_{\omega_{i}}}{1 - \alpha_{m_{i}+1} y_{n_{i}+1}}$$

$$+ \sum_{\substack{k=1\\k\neq\omega_{i}}}^{n_{i}} \pi_{\rm PH}(T_{k}^{i}(\omega)) \left(r_{k,\omega_{i}} + r_{k,n_{i}+1} \frac{\alpha_{m_{i}+1} y_{\omega_{i}}}{1 - \alpha_{m_{i}+1} y_{n_{i}+1}} \right).$$
(4.13)

The final balance equation concerns all states where node i is inactive, but at least one of its neighbors is active, so node i is blocked. The case of back-off freezing yields a trivial partial balance equation, as in this case no such state can be entered or exited due to changes in the state of node i. On the other hand, when the back-off process of blocked nodes is not frozen, the state can change due to a transition within the back-off process of node i:

$$\pi_{\mathrm{PH}}(\omega) \left(q_{\omega_{i},m_{i}+1} \left(1 - \frac{\alpha_{\omega_{i}}}{1 - \alpha_{m_{i}+1}} \right) + \sum_{\substack{k=1\\k\neq\omega_{i}}}^{m_{i}} q_{\omega_{i},k} \right)$$

$$= \sum_{\substack{k=1\\k\neq\omega_{i}}}^{m_{i}} \pi_{\mathrm{PH}}(T_{-k}^{i}(\omega)) \left(q_{k,\omega_{i}} + q_{k,m_{i}+1} \frac{\alpha_{\omega_{i}}}{1 - \alpha_{m_{i}+1}} \right), \quad \forall \omega \text{ s.t. } \omega_{i} \leq -1, \; \exists j \in C_{i} : \omega_{j} \geq 1.$$

$$(4.14)$$

We now proceed to show that π_{PH} from (4.1) indeed satisfies (4.12)-(4.14). Substituting π_{PH} into (4.12), and canceling common terms yields

$$\begin{split} &\beta_{\omega_i} \left(q_{\omega_i,m_i+1} \left(1 - \frac{\alpha_{\omega_i} \gamma_{n_i+1}}{1 - \alpha_{m_i+1} \gamma_{n_i+1}} \right) + \sum_{\substack{k=1\\k\neq\omega_i}}^{m_i} q_{\omega_i,k} \right) \\ &= \sum_{\substack{k=1\\k\neq\omega_i}}^{m_i} \beta_k \left(q_{k,\omega_i} + q_{k,m_i+1} \frac{\alpha_{\omega_i} \gamma_{n_i+1}}{1 - \alpha_{m_i+1} \gamma_{n_i+1}} \right) + \sigma_i \sum_{k=1}^{n_i} \eta_k r_{k,n_i+1} \frac{\alpha_{\omega_i}}{1 - \alpha_{m_i+1} \gamma_{n_i+1}}. \end{split}$$

By (4.11), and adding $\beta_{\omega_i} q_{\omega_i, m_i+1} \left(\frac{\alpha_{\omega_i} \gamma_{n_i+1}}{1-\alpha_{m_i+1} \gamma_{n_i+1}} - \frac{\alpha_{\omega_i}}{1-\alpha_{m_i+1}} \right)$ on both sides

$$\begin{split} \beta_{\omega_i} \left(q_{\omega_i,m_i+1} \left(1 - \frac{\alpha_{\omega_i}}{1 - \alpha_{m_i+1}} \right) + \sum_{\substack{k=1\\k\neq\omega_i}}^{m_i} q_{\omega_i,k} \right) &= \sum_{\substack{k=1\\k\neq\omega_i}}^{m_i} \beta_k q_{k,\omega_i} \\ &+ \sum_{k=1}^{m_i} \beta_k q_{k,m_i+1} \frac{\alpha_{\omega_i} \gamma_{n_i+1}}{1 - \alpha_{m_i+1} \gamma_{n_i+1}} + \nu_i \frac{\alpha_{\omega_i} (1 - \gamma_{n_i+1})}{1 - \alpha_{m_i+1} \gamma_{n_i+1}} - \beta_{\omega_i} q_{\omega_i,m_i+1} \frac{\alpha_{\omega_i}}{1 - \alpha_{m_i+1}}. \end{split}$$

4.A AUXILIARY RESULTS AND REMAINING PROOFS

By (4.10) and rearranging on both sides we get

$$\beta_{\omega_{i}}\left(q_{\omega_{i},m_{i}+1}\left(1-\frac{\alpha_{\omega_{i}}}{1-\alpha_{m_{i}+1}}\right)+\sum_{\substack{k=1\\k\neq\omega_{i}}}^{m_{i}}q_{\omega_{i},k}\right)=\sum_{\substack{k=1\\k\neq\omega_{i}}}^{m_{i}}\beta_{k}\left(q_{k,\omega_{i}}+q_{k,m_{i}+1}\frac{\alpha_{\omega_{i}}}{1-\alpha_{m_{i}+1}}\right)+\nu_{i}\frac{\alpha_{\omega_{i}}(1-\gamma_{n_{i}+1})}{1-\alpha_{m_{i}+1}\gamma_{n_{i}+1}}-\nu_{i}\frac{\alpha_{\omega_{i}}}{1-\alpha_{m_{i}+1}}(1-\alpha_{m_{i}+1}).$$

Canceling the remaining terms we get (4.8), so π_{PH} indeed satisfies (4.12). Substituting π_{PH} into (4.13), and canceling common terms yields

$$\sigma_{i}\eta_{\omega_{i}}\left(r_{\omega_{i},n_{i}+1}\left(1-\frac{\alpha_{m_{i}+1}y_{\omega_{i}}}{1-\alpha_{m_{i}+1}y_{n_{i}+1}}\right)+\sum_{\substack{k=1\\k\neq\omega_{i}}}^{n_{i}}r_{\omega_{i},k}\right)$$

$$=\sum_{k=1}^{m_{i}}\beta_{k}q_{k,m_{i}+1}\frac{y_{\omega_{i}}}{1-\alpha_{m_{i}+1}y_{n_{i}+1}}+\sigma_{i}\sum_{\substack{k=1\\k\neq\omega_{i}}}^{n_{i}}\eta_{k}\left(r_{k,\omega_{i}}+r_{k,n_{i}+1}\frac{\alpha_{m_{i}+1}y_{\omega_{i}}}{1-\alpha_{m_{i}+1}y_{n_{i}+1}}\right).$$

By (4.10), and adding $\sigma_i \eta_{\omega_i} r_{\omega_i, n_i+1} \left(\frac{\alpha_{m_i+1} \gamma_{\omega_i}}{1-\alpha_{m_i+1} \gamma_{n_i+1}} - \frac{\gamma_{\omega_i}}{1-\gamma_{n_i+1}} \right)$ on both sides

$$\sigma_{i}\eta_{\omega_{i}}\left(r_{\omega_{i},n_{i}+1}\left(1-\frac{\gamma_{\omega_{i}}}{1-\gamma_{n_{i}+1}}\right)+\sum_{\substack{k=1\\k\neq\omega_{i}}}^{n_{i}}r_{\omega_{i},k}\right)$$

$$=\nu_{i}\frac{\gamma_{\omega_{i}}(1-\alpha_{m_{i}+1})}{1-\alpha_{m_{i}+1}\gamma_{n_{i}+1}}+\sigma_{i}\sum_{\substack{k=1\\k\neq\omega_{i}}}^{n_{i}}\eta_{k}r_{k,\omega_{i}}+\sigma_{i}\frac{\alpha_{m_{i}+1}\gamma_{\omega_{i}}}{1-\alpha_{m_{i}+1}\gamma_{n_{i}+1}}\sum_{k=1}^{n_{i}}\eta_{k}r_{k,n_{i}+1}$$

$$-\sigma_{i}\eta_{\omega_{i}}r_{k,n_{i}+1}\frac{\gamma_{\omega_{i}}}{1-\gamma_{n_{i}+1}}.$$

By (4.11) and rearranging both sides we get

$$\sigma_{i}\eta_{\omega_{i}}\left(r_{\omega_{i},n_{i}+1}\left(1-\frac{\gamma_{\omega_{i}}}{1-\gamma_{n_{i}+1}}\right)+\sum_{\substack{k=1\\k\neq\omega_{i}}}^{n_{i}}r_{\omega_{i},k}\right)=\sigma_{i}\sum_{\substack{k=1\\k\neq\omega_{i}}}^{n_{i}}\eta_{k}\left(r_{k,\omega_{i}}+r_{k,n_{i}+1}\frac{\gamma_{\omega_{i}}}{1-\gamma_{n_{i}+1}}\right)$$
$$+\nu_{i}\frac{\gamma_{\omega_{i}}(1-\gamma_{n_{i}+1})}{1-\alpha_{m_{i}+1}\gamma_{n_{i}+1}}+\nu_{i}\frac{\alpha_{m_{i}+1}\gamma_{\omega_{i}}(1-\gamma_{n_{i}+1})}{1-\alpha_{m_{i}+1}\gamma_{n_{i}+1}}-\nu_{i}\frac{\gamma_{\omega_{i}}}{1-\gamma_{n_{i}+1}}(1-\gamma_{n_{i}+1}).$$

Canceling the remaining terms we get (4.9), so π_{PH} indeed satisfies (4.13).

Thirdly, π_{PH} trivially satisfies (4.14), as substitution of π_{PH} into this equation immediately gives (4.8).

STABILITY AND INSENSITIVITY

5 Fairness in linear networks

As mentioned in Section 1.3.2, a major drawback of CSMA-like protocols is unfairness, in the sense that some of the nodes get starved, while others receive high throughput. In the present chapter we study this unfairness in a linear network in which an active node blocks its neighbors on both sides. By choosing the back-off rate of each node as a particular function of the number of its neighbors, we can guarantee that all nodes in the network have the same throughput, completely removing the unfairness. We then investigate the consequences of this choice of activation rates on the network-average saturated throughput, and we show that these rates perform well in non-saturated settings.

Although we assume that back-off times and transmission durations are exponentially distributed, we know by Theorem 4.1 that all results hold for generally distributed back-off times and transmission durations as well. In Chapter 6 we provide an alternative proof for the main result in this chapter, using Markov random fields.

This chapter is structured as follows. In Section 5.1 we introduce the linear network in more detail. In Section 5.2 we study some of the key features of the unfairness that arises when all nodes have equal back-off rates. In Section 5.3 we determine the fair back-off rates that yield equal throughputs. In Section 5.4 we investigate the impact of the fair back-off rates on the network-average throughput and in Section 5.5 we discuss the performance of the fair back-off rates in an unsaturated network. Section 5.6 presents some conclusions.

FAIRNESS IN LINEAR NETWORKS

5.1 Model description

We model the unfairness using the CSMA model introduced in Chapter 1. We consider a linear network of *n* nodes on a line, where a transmitting node blocks the first β nodes on both sides. So the conflict graph is such that all vertices $|i - j| \le \beta$ are connected, and the set of feasible states Ω is given by all ω such that no two 1's in ω are β positions or less apart, i.e., $\omega_i \omega_i = 0$ if $1 \le |i - j| \le \beta$.

Alternatively, we can express the set of feasible states as all states that satisfy a certain system of linear equations. Let *A* be an $(n - \beta) \times n$ matrix where each row contains $\beta + 1$ consecutive 1's, in the following way:

$$A = \begin{pmatrix} 1 & 1 & \dots & 1 & 0 & \dots & 0 & 0 \\ 0 & 1 & 1 & \dots & 1 & 0 & \dots & 0 \\ & & \ddots & & & \ddots & & \vdots \\ 0 & \dots & 0 & 1 & 1 & \dots & 1 & 0 \\ 0 & 0 & \dots & 0 & 1 & 1 & \dots & 1 \end{pmatrix}.$$
 (5.1)

Now we can write the state space as $\Omega = \{\omega \in \{0,1\}^n \mid A\omega \leq 1\}$, where **1** is the all-1 vector (of size $n - \beta$). This characterization has a natural interpretation as a set of capacity constraints, and nodes can activate only when enough capacity is available. We allocate unit capacity to each node, and use the convention that whenever a node is active it uses its own capacity, as well as the capacity of all its neighbors to the left. The *i*th row of *A* thus represents the capacity required when node *i* is active. The constraints that arise from the first β nodes on the line are redundant, and ignoring these leads to the matrix *A* in (5.1).

We assume that all nodes are saturated and that unblocked nodes activate after an exponentially distributed (back-off) time with mean $1/\nu_i$. Without loss of generality, we assume that transmissions last for an exponentially distributed time with unit mean. Under these assumptions, the *n*-dimensional process that describes the node activity is a continuous-time Markov process. We have seen that the stationary distribution of this process is given by (see (1.8))

$$\pi(\omega) = \begin{cases} Z_n^{-1} \prod_{i=1}^n v_i^{\omega_i}, & \text{if } \omega \in \Omega, \\ 0, & \text{otherwise,} \end{cases}$$
(5.2)

where Z_n is the normalization constant of an *n*-node linear network. Note that we made the dependence of the normalization constant on the network length explicit, which we will use elsewhere in this chapter. From Chapter 4 we know that the distribution (5.2) also holds for generally distributed back-off times and transmission durations. Since all results in this chapter are based on (5.2), they remain valid for general distributions.

Our main concern is with the long-term behavior of nodes, characterized by their throughputs. We study the throughput vector $\theta = (\theta_1, ..., \theta_n)$, where θ_i represents the fraction of time node *i* is active. Recall that

$$\boldsymbol{\theta} = \sum_{\boldsymbol{\omega} \in \Omega} \boldsymbol{\pi}(\boldsymbol{\omega}) \boldsymbol{\omega}.$$
 (5.3)

By exploiting the structure of the network, we can construct alternative expressions for the throughput in (5.3). More specifically, we make use of the observation that if node *i* is active, nodes to the left of *i* behave independently from nodes to the right of *i*. This leads to the following theorem.

THEOREM 5.1. Define the sequence $(Z_i)_{i=-\infty}^{\infty}$ such that $Z_i = 1$ for $i \leq 0$, and

$$Z_i = 1 + \nu_1 + \dots + \nu_i, \quad i = 1, 2, \dots, \beta + 1,$$
(5.4)

$$Z_i = Z_{i-1} + \nu_i Z_{i-\beta-1}, \quad i = \beta + 2, \beta + 3, \dots$$
(5.5)

Let the vector of back-off rates $\nu = (v_1, ..., v_n)$ be such that $v_i = v_{n+1-i}$, i = 1, ..., n. Then

$$\theta_i = \nu_i \frac{Z_{i-\beta-1} Z_{n-i-\beta}}{Z_n}, \quad i = 1, ..., n.$$
(5.6)

PROOF. By conditioning on whether or not node i is active, we can decompose the activity of the network into two parts, separated by this active node (see [11, Equation (15)]),

$$\theta_i = \nu_i \frac{Z_{1:i-\beta-1} Z_{i+\beta+1:n}}{Z_{1:n}},$$
(5.7)

where $Z_{i:j}$ is the normalization constant of a network consisting only of nodes i, \ldots, j . For simplicity we denote $Z_i = Z_{1:i}$, and the symmetry of ν implies

$$Z_{i:n} = Z_{1:n-i+1}. (5.8)$$

Substituting (5.8) into (5.7) yields the expression for θ_i in (5.6). By conditioning on the activity of node *i*, we immediately get the recursion relation (5.5).

5.2 Unfairness

We now venture deeper into the problem of unfairness, and assume for now that all nodes have equal back-off rates $v_i = \sigma$. As observed in Section 1.4 the throughput distribution in this case is highly unfair, in the sense that some nodes have a larger throughput than others. In this section we evaluate the throughput in order to study the unfairness in more detail.

In order to compute the throughput in (5.6), we need to compute the Z_i from (5.4) and (5.5). A detailed analysis of the Z_i is performed in Chapter 7, where it is shown that

$$Z_{i} = \sum_{j=0}^{\beta} c_{j} \lambda_{j}^{i}, \quad i = 0, 1, \dots,$$
(5.9)

with

$$c_j = \frac{\lambda_j^{p+1}}{(\beta+1)\lambda_j - \beta},\tag{5.10}$$

and $\lambda_0, \ldots, \lambda_\beta$ the $\beta + 1$ roots of

$$\lambda^{\beta+1} - \lambda^{\beta} - \sigma = 0. \tag{5.11}$$

Moreover, with λ_0 representing the root with the largest modulus of (5.11), it follows from (5.9) that

$$Z_i = c_0 \lambda_0^i (1 + o(1)), \quad i \to \infty,$$
 (5.12)



Figure 5.1: The per-node throughput for $\beta = 1$ and various values of *n* and σ .

and we will use this asymptotic relation at several places, see (5.13) and (5.45).

For ease of presentation, we restrict ourselves to $\beta = 1$ in the remainder of this section. Figures 5.1(a)-5.1(d) show the per-node throughput for various values of *n* and σ . All figures display a similar pattern, with the outer nodes having the highest throughput. Moreover, all figures are symmetric, and exhibit some form of oscillatory behavior. These observations are formalized in the following result.

PROPOSITION 5.1. For $v_i = \sigma > 0$, i = 1, ..., n and $\beta = 1$, the throughput has the following properties:

(i) Symmetric: θ_i = θ_{n-i+1}, i = 1, 2, ..., n.
(ii) Alternating and converging: (-1)ⁱ(θ_{i+1} - θ_i) is positive and decreasing for i = 1, 2, ..., [n/2].

Proposition 5.1 is proved in Appendix 5.A.

In Figure 5.1 we see that for $\beta = 1$, the largest difference in throughput is between nodes 1 and 2. Proposition 5.1(ii) confirms that this is the most unfair situation, and it persists even in larger networks where the node-in-the-middle problem is mitigated by the activity of the remaining nodes. In fact, for large networks we have the following result.

PROPOSITION 5.2. For $v_i = \sigma > 0$, i = 1, ..., n and $\beta = 1$,

$$\frac{\theta_1}{\theta_2} \sim \lambda_0 = \frac{1 + \sqrt{1 + 4\sigma}}{2}, \quad n \to \infty.$$

5.3 ACHIEVING FAIRNESS

PROOF. We have from (5.6) that $\theta_1 = Z_{n-2}/Z_n$ and $\theta_2 = Z_{n-3}/Z_n$. Using (5.12) we obtain

$$\frac{\theta_1}{\theta_2} \sim \lambda_0, \quad n \to \infty.$$
 (5.13)

For $\beta = 1$ we can explicitly solve (5.11) to obtain $\theta_1/\theta_2 \sim \lambda_0 = \frac{1}{2}(1 + \sqrt{1 + 4\sigma})$.

We note that for $\beta = 1$, the Z_i satisfy a three-term recursion reminiscent of that satisfied by the Chebyshev polynomials U_n of the second kind. Accordingly, we have

$$Z_{i} = (-\sigma)^{\frac{1}{2}(i+1)} U_{i+1}(\sqrt{-1/4\sigma}) = \sum_{j=0}^{\lfloor \frac{i+1}{2} \rfloor} {i+1-j \choose j} \sigma^{j}.$$

The latter expression can be interpreted as the summation over all possible combinations of nodes that can be active simultaneously.

Results similar to those presented in this section can be obtained for $\beta \ge 2$. As an example, Figures 5.2(a)-5.2(b) show the per-node throughput for n = 9 and $\beta = 2, 3$. Both figures exhibit similar oscillatory behavior as observed for $\beta = 1$, although the oscillation period increases with β .



Figure 5.2: The per-node throughput for n = 9 and various values of β and σ .

5.3 Achieving fairness

In this section we present a way to completely remove the unfairness that was discussed in Section 5.2. In order to do so, we choose node-dependent back-off rates v_i such that all nodes have equal throughputs ($\theta_1 = \theta_2 = \cdots = \theta_n$). From (5.2) and (5.3) we see that in order to meet this objective we have to solve a system of *n* nonlinear equations. It seems that in general this system cannot be solved directly. We therefore choose a more indirect approach, and we first consider two special cases that can be solved explicitly. The insight obtained from these exact solutions is then used to guess the general solution to the system of nonlinear equations.

The first case is when $\beta = n - 2$, so that all but the two outer nodes will block the entire network.

PROPOSITION 5.3. For linear networks with three or more nodes, and $\beta = n - 2$, setting $v_1 = v_n = \alpha$, $v_i = \alpha(1 + \alpha)$ for all other nodes and $\alpha > 0$, yields equal throughputs

$$\theta_i = \frac{\alpha}{1 + (\beta + 1)\alpha}, \quad i = 1, \dots, n.$$
(5.14)

PROOF. From (5.6) we see that

$$\theta_1 = Z_n^{-1} \nu_1 (1 + \nu_n), \tag{5.15}$$

$$\theta_i = Z_n^{-1} \nu_i, \quad i = 2, 3, \dots, n-1,$$
(5.16)

$$\theta_n = Z_n^{-1} \nu_n (1 + \nu_1). \tag{5.17}$$

The inherent symmetry of the model allows us to set $v_1 = v_n$. Moreover, for the throughput of the other nodes to be equal, we require $v_2 = \cdots = v_{n-1} = v_1(1 + v_1)$. If we set $v_1 = \alpha$, and substitute this into (5.15)-(5.17), we get a throughput of

$$\theta_i = Z_n^{-1} \alpha (1 + \alpha). \tag{5.18}$$

The normalization constant Z_n can be determined by summing over all feasible states:

$$Z_n = 1 + \sum_{i=1}^n v_i + v_1 v_n = 1 + (n-2)\alpha(1+\alpha) + 2\alpha + \alpha^2$$

= (1 + \alpha)(1 + (\beta + 1)\alpha). (5.19)

Substituting (5.19) into (5.18) yields (5.14).

The case n = 5, $\beta = 3$ of Proposition 5.3 was considered in [17]. The second special case corresponds to $n = 2(\beta + 1)$, so that a node blocks at least half of the network.

PROPOSITION 5.4. For linear networks with n = 2m nodes, $m \in \mathbb{N}$, and $\beta = m - 1$, setting $v_i = \alpha(1 + \alpha)^{i-1}$ for $\alpha > 0$ and i = 1, ..., m yields equal throughputs

$$\theta_i = \frac{\alpha}{1 + (\beta + 1)\alpha}, \quad i = 1, \dots, n.$$
(5.20)

PROOF. To achieve equal throughputs, we see from (5.2) and (5.3) that for the case at hand we should solve the system of equations

 $v_1 + v_1(v_{m+1} + \dots + v_n) = v_2 + v_2(v_{m+2} + \dots + v_n) = \dots = v_m + v_m v_n.$ (5.21)

Indeed, the throughput of node i can be written as a sum over all states in which node i is active. Using symmetry, (5.21) can be written as

$$v_1 + v_1(v_1 + \dots + v_m) = v_2 + v_2(v_1 + \dots + v_{m-1}) = \dots = v_m + v_m v_1.$$
 (5.22)

Let $v_1 = \alpha > 0$. The solution of (5.22) is easily seen to be $v_i = \alpha(1+\alpha)^{i-1}$, i = 1, ..., m, and hence

$$\theta_i = Z_n^{-1} \alpha (1+\alpha)^m. \tag{5.23}$$

Summing over all possible states yields

$$Z_n = 1 + \sum_{i=1}^n \nu_i + \sum_{i=1}^m \nu_i \sum_{j=i+m}^n \nu_j = (1 + (\beta + 1)\alpha)(1 + \alpha)^m.$$
(5.24)

Substituting (5.24) into (5.23) gives (5.20).

5.3 ACHIEVING FAIRNESS

It is clear that the complexity of the system of equations governed by (5.3) reduces considerably for the choices of β discussed in Propositions 5.3 and 5.4. For general β this system remains rather complicated. However, we can use Propositions 5.3 and 5.4 to make an educated guess for the general solution. First observe that the fair back-off rates in Propositions 5.3 and 5.4 only depend on the number of neighbors (nodes within β hops). Denote by $\gamma(i)$ the number of neighbors of node *i*, let $\alpha > 0$ be some positive constant, and choose back-off rates ν_i^* as

$$\nu_i^* = \alpha (1+\alpha)^{\gamma(i)-\gamma(1)}.$$
 (5.25)

We see that this choice is consistent with the fair back-off rates in Propositions 5.3 and 5.4. We now show that v_i^* indeed achieves fairness for all β .

THEOREM 5.2. Let $\alpha > 0$, $\beta \le n - 1$ and choose v_i^* as in (5.25). Then

$$\theta_i = \frac{\alpha}{1 + (1 + \beta)\alpha}, \quad i = 1, \dots, n.$$
(5.26)

We first show that when the back-off rates are chosen according to (5.25), the recursive relations (5.4) and (5.5) for the normalization constant Z_i have a closed-form solution.

LEMMA 5.1. Let $\alpha > 0$ and choose v_i^* as in (5.25). Then

$$Z_i = (1 + \alpha)^i, \quad i = 1, 2, \dots, n - \beta.$$
(5.27)

PROOF. Substituting (5.27) into (5.4) gives,

$$Z_i = 1 + \alpha + \alpha(1 + \alpha) + \cdots + \alpha(1 + \alpha)^{i-1} = (1 + \alpha)^i,$$

for $i \le \beta + 1$. Substituting (5.27) into (5.5) gives,

$$Z_i = (1 + \alpha)^{i-1} + \alpha (1 + \alpha)^{\beta} (1 + \alpha)^{i-\beta-1} = (1 + \alpha)^i,$$

for $i \ge \beta + 2$.

With Lemma 5.1 we are now in position to prove our main result.

PROOF OF THEOREM 5.2. Recall from (5.6) that

$$\theta_i = \nu_i \frac{Z_{i-\beta-1} Z_{n-i-\beta}}{Z_n}, \quad i = 1, ..., n.$$
(5.28)

To prove Theorem 5.2 we substitute (5.27) into (5.28). We distinguish between different values of i.

For $i \ge \beta + 1$ and $i \le n - \beta$ we see that $v_i^* = \alpha (1 + \alpha)^{\beta}$ and

$$Z_{i-\beta-1} = (1+\alpha)^{i-\beta-1}, \quad Z_{n-i-\beta} = (1+\alpha)^{n-i-\beta}.$$
 (5.29)

Similarly, for $i \ge \beta + 1$ and $i \ge n - \beta + 1$ we have $v_i^* = \alpha(1 + \alpha)^{n-i}$ and

$$Z_{i-\beta-1} = (1+\alpha)^{i-\beta-1}, \quad Z_{n-i-\beta} = 1.$$
 (5.30)

FAIRNESS IN LINEAR NETWORKS

For $i \leq \beta$ and $i \leq n - \beta$ we have $v_i^* = \alpha (1 + \alpha)^{i-1}$ and

$$Z_{i-\beta-1} = 1, \quad Z_{n-i-\beta} = (1+\alpha)^{n-i-\beta}.$$
 (5.31)

Finally, for $i \leq \beta$ and $i \geq n - \beta + 1$ we have $v_i^* = \alpha (1 + \alpha)^{n - \beta - 1}$ and

$$Z_{i-\beta-1} = 1, \quad Z_{n-i-\beta} = 1.$$
 (5.32)

Substituting (5.29)-(5.32) into (5.28) yields

$$\theta_i = Z_n^{-1} \alpha (1+\alpha)^{n-\beta-1}.$$
 (5.33)

We next consider the normalization constant. With *m* such that $n = \beta + m$, by (5.5),

$$Z_n = Z_{n-1} + \nu_n^* Z_{n-\beta-1},$$

which gives upon iteration

$$Z_n = Z_{n-\beta} + \sum_{i=1}^{\beta} v_{n+1-i}^* Z_{n-\beta-i}.$$
(5.34)

Substituting (5.27) into (5.34) yields

$$Z_{n} = (1+\alpha)^{n-\beta} + \sum_{i=1}^{\min\{m,\beta\}} \alpha (1+\alpha)^{i-1} (1+\alpha)^{n-\beta-i} + \sum_{i=m+1}^{\beta} \alpha (1+\alpha)^{n-\beta-i}$$
$$= (1+\alpha)^{n-\beta-1} (1+(\beta+1)\alpha).$$
(5.35)

Combining (5.35) and (5.33) leads to (5.26).

In Chapter 6 we provide an alternative proof of Theorem 5.2, using Markov random fields.

To understand better why the rates (5.25) only depend on the number of neighbors of each node, we study the rates in the limiting regimes of light traffic ($\alpha \downarrow 0$) and heavy traffic ($\alpha \rightarrow \infty$). First write (5.25) as

$$\nu_i^* = \alpha \sum_{j=0}^{\gamma(i) - \gamma(1)} {\gamma(i) - \gamma(1) \choose j} \alpha^j, \quad i = 1, \dots, n.$$
(5.36)

When α is small, nodes activate slowly, and few nodes will be active simultaneously. In fact, the Markov process spends most of its time in states with at most one active node, and node interaction (blocking) is negligible. This is reflected in the light-traffic back-off rates that follow immediately from (5.36):

$$v_i^* = \alpha + (\gamma(i) - \gamma(1))\alpha^2 + O(\alpha^3), \quad \alpha \downarrow 0.$$

Hence, for small α , $v_i^* \approx \alpha$, which is the same for all nodes. Indeed, if at most one node is active (as is the case for α small), there is no blocking, and therefore no need to discriminate between nodes. As α increases, states with two active nodes are increasingly likely, and nodes may now block their neighbors (all nodes within distance β). This is accounted for in the back-off rate by the term $(\gamma(i) - \gamma(1))\alpha^2$,

5.4 NETWORK-AVERAGE THROUGHPUT

which is linear in the number of neighbors. Thus in light traffic, only the number of neighbors is of importance, rather than the structure of the entire network. This reasoning extends to more general networks.

Next, we consider large back-off rates, and we compare the equal rates (all nodes activate with rate $\sigma \rightarrow \infty$) with the fair rates (5.25) (with $\alpha \rightarrow \infty$). In both cases, nodes activate almost instantaneously when they get the chance to do so, i.e., when all neighbors are inactive. Consequently, the only states that have positive probability in the limit are those consisting of maximal independent sets of active nodes. The distribution according to which these maximal states occur depends on the choice of back-off rates.

First consider the case of equal back-off rates $v_i = \sigma$, i = 1, ..., n. We have seen in Section 5.2 that this creates unfairness, and that the unfairness increases with σ . In particular, we see from (5.2) that the only states that have positive probability for $\sigma \rightarrow \infty$ are those of maximum size, i.e., states with $\lceil n/(\beta + 1) \rceil$ active nodes. Thus, for $\sigma \rightarrow \infty$,

$$\pi(\boldsymbol{\omega}) = \begin{cases} 1/|\mathcal{M}|, & \text{if } \boldsymbol{\omega} \in \mathcal{M}, \\ 0, & \text{otherwise,} \end{cases}$$

with $\mathcal{M} \subset \Omega$ the set of states of maximum size and $|\mathcal{M}|$ the cardinality of this set. The throughput of each node is thus determined by the number of maximum states it is contained in, which is not necessarily the same for all nodes.

For the fair back-off rates (5.36), we see that

$$\nu_i^* = \alpha^{\gamma(i) - \gamma(1) + 1} + O(\alpha^{\gamma(i) - \gamma(1)}), \quad \alpha \to \infty.$$
(5.37)

Thus the back-off rate of a node is characterized by the leading exponent y(i) - y(1) + 1, and the limiting probability of a state is determined by the sum of these exponents over all active nodes. In fact, the only states that give a contribution for $\alpha \to \infty$ are those that maximize the sum of the exponents of α over all active nodes. It turns out that there are $\beta + 1$ such states, with active nodes $\{i, i + (\beta + 1), \dots, i + (\beta + 1)\kappa_i\}$, $i = 1, \dots, \beta + 1$, with $\kappa_i = \lceil \frac{n+1-i}{\beta+1} \rceil - 1$. Each such state has limiting probability

$$\pi(\omega)=Z_n^{-1}\prod_{j=0}^{\kappa_i}\nu_{i+(\beta+1)j}=Z_n^{-1}(\alpha^{n-\beta}+O(\alpha^{n-\beta-1}))=\frac{1}{\beta+1},\quad \alpha\to\infty,$$

because $Z_n = (\beta + 1)\alpha^{n-\beta} + O(\alpha^{n-\beta-1})$, as $\alpha \to \infty$. Contrary to the case of equal rates, we see that each node appears in exactly one state with positive limiting probability. This explains the equal throughputs in heavy traffic.

This result strongly depends on the structure of the network, as the maximal independent sets may change drastically with the addition or the removal of even a single node. As a result, the simple, locally determined heavy-traffic back-off rates (5.37) may only hold for linear networks.

5.4 Network-average throughput

The fair rates v_i^* in (5.25) are designed to remove the unfairness that arises when all nodes have equal back-off rates σ . In order to compare the two schemes, we want to

set their respective parameters α and σ such that the average per-node throughputs are equal. In a network with $v_i = \sigma > 0$, i = 1, ..., n, write $Z_i(\sigma)$ and $\theta_i(\sigma)$ for the normalization constant of a network with *i* nodes, and the throughput of node *i*, respectively. Let $\bar{\theta}_n(\sigma) = n^{-1} \sum_{i=1}^n \theta_i(\sigma)$ denote the average per-node throughput in a network with *n* nodes.

In Section 5.3 we showed that all nodes have equal throughputs $\alpha/(1 + \alpha(\beta + 1))$ when using the fair back-off rates in (5.25). When all nodes have equal back-off rates, a closed-form expression for the throughput does not seem available. However, we can express the average per-node throughput in terms of the normalization constant Z_n .

PROPOSITION 5.5. Let $v_i = \sigma$, i = 1, ..., n. The average per-node throughput is given by

$$\bar{\theta}_n(\sigma) = \frac{\sigma}{nZ_n(\sigma)} \frac{\mathrm{d}Z_n(\sigma)}{\mathrm{d}\sigma}.$$
(5.38)

PROOF. We have from (5.6) with $v_i = \sigma$ that

$$\bar{\theta}_n(\sigma) = \frac{\sigma}{nZ_n} \sum_{i=1}^n Z_{i-\beta-1} Z_{n-i-\beta}$$

We compute, using the definition of Z_i in Theorem 5.1,

$$\sum_{n=1}^{\infty} \Big(\sum_{i=1}^{n} Z_{i-\beta-1} Z_{n-i-\beta} \Big) x^n = x \Big(\frac{x^{\beta}-1}{x-1} + x^{\beta} G_Z(x) \Big)^2,$$

with $G_Z(x)$ the generating function of the Z_i given by

$$G_Z(x) = \sum_{i=0}^{\infty} Z_i x^i = \frac{x - 1 + \sigma x^{\beta+1} - \sigma x}{(x - 1)(1 - x - \sigma x^{\beta+1})},$$

see (7.2). Some rewriting then gives

$$\sum_{n=1}^{\infty} \left(\sum_{i=1}^{n} Z_{i-\beta-1} Z_{n-i-\beta} \right) x^n = \frac{\chi}{(1-\chi-\sigma x^{\beta+1})^2}$$

On the other hand, we compute that

$$\frac{\mathrm{d}}{\mathrm{d}\sigma} [G_Z(x)] = \frac{x}{(1 - x - \sigma x^{\beta + 1})^2}$$

and the result follows.

By Proposition 5.5 and the expression (5.9) for Z_i , we can express the average per-node throughput in terms of the roots $\lambda_0, \ldots, \lambda_\beta$ of (5.11).

PROPOSITION 5.6. Let $v_i = \sigma$, i = 1, ..., n. The average per-node throughput is given by

$$\bar{\theta}_n(\sigma) = \frac{\sigma}{n} \frac{P}{Q},\tag{5.39}$$

where

$$P = \sum_{j=0}^{\beta} \frac{\lambda_j^{n+1}}{(\beta+1)\lambda_j - \beta} \left(\frac{n+\beta+1}{(\beta+1)\lambda_j - \beta} - \frac{(\beta+1)\lambda_j}{((\beta+1)\lambda_j - \beta)^2} \right),$$
$$Q = \sum_{j=0}^{\beta} \frac{\lambda_j^{n+\beta+1}}{(\beta+1)\lambda_j - \beta}.$$
(5.40)

PROOF. By (5.9) and (5.10) we have

$$Z_n(\sigma) = \sum_{j=0}^{\beta} \frac{\lambda_j^{n+\beta+1}}{(\beta+1)\lambda_j - \beta},$$
(5.41)

where λ_j are the $(\beta + 1)$ roots λ of (5.11). By implicit differentiation of (5.11) with respect to σ we find

$$\frac{\mathrm{d}\lambda_j}{\mathrm{d}\sigma} = \frac{1}{\lambda_j^{\beta-1}} \frac{1}{(\beta+1)\lambda_j - \beta}.$$
(5.42)

Then from (5.41) and (5.42) we get

$$\frac{\mathrm{d}Z_n(\sigma)}{\mathrm{d}\sigma} = \sum_{j=0}^{\beta} \left(\frac{(n+\beta+1)\lambda_j^{n+\beta}}{(\beta+1)\lambda_j - \beta} - \frac{(\beta+1)\lambda_j^{n+\beta+1}}{((\beta+1)\lambda_j - \beta)^2} \right) \frac{\mathrm{d}\lambda_j}{\mathrm{d}\sigma}$$
$$= \sum_{j=0}^{\beta} \frac{\lambda_j^{n+1}}{(\beta+1)\lambda_j - \beta} \left(\frac{n+\beta+1}{(\beta+1)\lambda_j - \beta} - \frac{(\beta+1)\lambda_j}{((\beta+1)\lambda_j - \beta)^2} \right). \tag{5.43}$$

The result follows from substituting (5.41) and (5.43) into (5.38).

When the network grows large $(n \rightarrow \infty)$ the root of largest modulus, λ_0 , becomes dominant, and (5.39)-(5.40) simplifies.

COROLLARY 5.1. Let $v_i = \sigma$, i = 1, ..., n. The limiting average per-node throughput $\bar{\theta}(\sigma) = \lim_{n \to \infty} \bar{\theta}_n(\sigma)$ is given by

$$\bar{\theta}(\sigma) = \frac{\lambda_0 - 1}{(\beta + 1)\lambda_0 - \beta}.$$
(5.44)

PROOF. We have, as $n \to \infty$,

$$P = \frac{\lambda_0^{n+1}}{(\beta+1)\lambda_0 - \beta} \frac{n}{(\beta+1)\lambda_0 - \beta} (1 + o(1)), \quad Q = \frac{\lambda_0^{n+\beta+1}}{(\beta+1)\lambda_0 - \beta} (1 + o(1)).$$
(5.45)

Hence

$$\bar{\theta}_n(\sigma) = \frac{\sigma}{n} \frac{P}{Q} = \frac{\sigma \lambda_0^{-\beta}}{(\beta+1)\lambda_0 - \beta} (1 + o(1)),$$

and the result follows as $\sigma \lambda_0^{-\beta} = \lambda_0 - 1$ by (5.11).

The limiting expression (5.44) occurs in a variety of contexts in [6, 22, 71, 93, 106]. When $\beta \sigma \rightarrow \infty$, the throughput (5.44) simplifies even further.

79

FAIRNESS IN LINEAR NETWORKS

COROLLARY 5.2. Let $v_i = \sigma$, i = 1, ..., n and let $n \to \infty$. The limiting average throughput satisfies

$$\bar{\theta}(\sigma) = \frac{1}{\beta+1}(1+o(1)), \quad \beta\sigma \to \infty.$$

PROOF. By rewriting (5.44) we have

$$\bar{\theta}(\sigma) = \frac{1}{\beta + 1} \frac{1}{1 + \frac{1}{(\beta + 1)(\lambda_0 - 1)}}.$$

Consequently, for $\bar{\theta}(\sigma) = \frac{1}{\beta+1}(1 + o(1))$ to hold, it is necessary and sufficient that $(\beta + 1)(\lambda_0 - 1) \rightarrow \infty$. Recall from (5.11) that λ_0 is such that

$$\lambda_0^p (\lambda_0 - 1) = \sigma. \tag{5.46}$$

Let M > 0 be some positive constant, and assume that $\beta \sigma \leq M$. Then

$$\beta \lambda_0^{\beta} (\lambda_0 - 1) = \beta \sigma \le M,$$

and so $\beta(\lambda_0 - 1) \le M$. Conversely, assume that $\beta(\lambda_0 - 1) \le K$ for some positive constant K > 0. Then

$$\beta\sigma = \beta(\lambda_0 - 1)\lambda_0^\beta \le \beta(\lambda_0 - 1)\exp(\beta(\lambda_0 - 1)) \le Ke^K.$$

Hence

$$\beta(\lambda_0 - 1)$$
 bounded $\Leftrightarrow \beta \sigma$ bounded.

It follows that a sufficient condition for $(\beta + 1)(\lambda_0 - 1) \rightarrow \infty$ is that $\beta \sigma \rightarrow \infty$.

Corollary 5.2 implies that $\bar{\theta}(\sigma) \rightarrow \frac{1}{\beta+1}$ for β fixed and $\sigma \rightarrow \infty$. Thus for $n \rightarrow \infty$, both the equal and fair back-off rates can achieve the maximum throughput by letting $\sigma \rightarrow \infty$ and $\alpha \rightarrow \infty$, respectively.

Next, we fix $\sigma > 0$ and search for $\alpha = \alpha_n(\sigma)$ such that

$$\bar{\theta}_n(\sigma) = \frac{\alpha}{1 + \alpha(\beta + 1)},\tag{5.47}$$

so the network-average throughput is identical for the fair rates and equal rates. For $\alpha(\sigma) = \lim_{n\to\infty} \alpha_n(\sigma)$ we can make this comparison explicit. By equating (5.26) and (5.44) and solving for α , we have $\alpha(\sigma) = \lambda_0 - 1$.

It is intuitively clear that imposing fairness may compromise the throughput. From (5.26) it is seen that the fair per-node throughputs are bounded from above by $\frac{1}{\beta+1}$, and that this upper bound can be approached arbitrarily closely by letting $\alpha \to \infty$. Corollary 5.1 shows that, as $n \to \infty$, the average throughputs in the fair case and unfair case are equal when α is taken to be $\lambda_0 - 1$. The maximum back-off rate in this limiting case equals

$$v_{\max}^* = \alpha (1+\alpha)^{\beta} = v_i, \quad \beta + 1 \le i \le n - \beta$$

as is seen from (5.25). This maximum equals σ by (5.46) since $1 + \alpha = \lambda_0$. Hence, as $n \to \infty$, the fair case achieves the same average throughput with back-off rates that are less than or equal to those in the unfair case.



5.5 Fair back-off rates in unsaturated networks

Throughout this chapter we have assumed all nodes to be saturated, and we derived fair back-off rates that give equal throughputs to all nodes. Alternatively, we may consider a model where packets arrive over time, and nodes may not always have packets to transmit. An active node transmits one packet, which takes an exponential time with mean 1. After finishing this transmission, the node has to compete for access again. The dynamics are the same as in the saturated scenario, except that when a node has no packets for transmission, it will remain inactive until it receives a new packet. We argue that the fair back-off rates also perform well in such unsaturated settings.

81



Figure 5.3: Two types of unsaturated networks.

First consider the situation in Figure 5.3(a), where packets arrive at each node according to an independent Poisson process with rate λ , and leave the system once transmitted. Nodes activate according to the fair back-off rates (5.25) with parameter $\alpha > 0$. This model reduces to the saturated model when $\lambda \to \infty$. Simulations suggest that the system is stable whenever $\lambda < \alpha/(1 + \alpha(\beta + 1))$, the saturation throughput.

Next, consider the situation in Figure 5.3(b), where packets arrive at node 1 according to a Poisson process with rate λ , and are routed along nodes 2,..., n. When node n finishes a transmission, the corresponding packet exits the system. The throughput of node n is of special interest, as it represents the end-to-end throughput of the network, that is, the rate at which packets leave the network. If $\theta_n = \lambda$, the system will eventually empty. If $\theta_n < \lambda$, packets arrive at a higher rate than the network can sustain, and packets will accumulate at certain bottleneck nodes. Figure 5.4 shows simulation results for the end-to-end throughput of this network, plotted against the arrival rate λ for n = 5, $\beta = 1$. The thick, gray line corresponds to the network where all nodes have equal back-off rate $\sigma = 6$, and the black line shows the throughput of a network with fair back-off rates (5.25) and $\alpha = 11.68$. The values of σ and α are chosen such that the average per-node throughput is equal for both back-off schemes, as prescribed by (5.47). The network with equal back-off rates performs poorly. When the arrival rate grows beyond a certain threshold, node 2 saturates and the throughput drops [18, 102]. The network with fair rates, on the other hand, can sustain higher arrival rates and experiences no decrease in throughput when in overload. In fact, the end-to-end throughput approaches the per-node throughput in the corresponding saturated network (indicated by the dashed horizontal line). So the network again appears to be stable when $\lambda < \alpha/(1 + \alpha(\beta + 1))$.





Figure 5.4: The end-to-end throughput of a network with equal back-off rates (thick, gray) and fair back-off rates (black), plotted against the arrival rate at node 1.

5.6 Concluding remarks

In this chapter we studied unfairness and fairness in linear CSMA networks. We proposed node-specific fair back-off rates (5.25) as a function of the number of neighbors, and showed that these rates provide equal throughput for all nodes. The fair back-off rates increase with the number of neighbors. Intuitively, this can be explained by the observation that highly contended nodes require higher back-off rates to remain competitive. Consequently, the rates (5.25), which are exact in linear networks, might serve as a heuristic in more complex networks.

These node-dependent back-off rates are still in line with the distributed nature of CSMA protocol, as calculating the fair back-off rates only requires the number of neighbors, which a node can observe locally by sensing its direct environment. Finding exact expressions for the back-off rates that provide strict fairness for networks beyond the linear network is challenging. In [39] results were obtained for trees with $\beta = 1$, and it was shown that rates such as in (5.25), where nodes on the leaves of the tree have lower rates than those in the stem of the tree, provide strict fairness. For such trees, it seems possible to extend this result to the β -hop blocking situation. Other regular topologies such as certain grids appear to admit a similar analysis as well.

The results obtained in this chapter rely heavily on the decomposition (5.6), which only applies for certain well-structured networks. For more general networks, the objective of equal throughputs boils down to solving the system of nonlinear equations that follows from (5.3). This problem is addressed in Chapter 6 for general conflict graphs. There we will show that the throughput function that maps the back-off rates to the throughput is in fact globally invertible, meaning that for any feasible throughput vector (fair or otherwise) there exists a unique vector of back-off rates that yields this throughput vector. We then propose various numerical algorithms to determine the appropriate back-off rates. 5.A PROOF OF PROPOSITION 5.1

Appendix

5.A Proof of Proposition 5.1

We first establish an auxiliary result. Define a(i, l, n) as the number of states in which exactly l nodes are active, including node i. For successive nodes, the following relations hold.

LEMMA 5.2. For $n \in \mathbb{N}$, $i \leq \left\lceil \frac{n}{2} \right\rceil - 1$,

$$a(i,l,n) = a(i+1,l,n), \quad l \le i,$$
 (5.48)

83

$$a(i,l,n) > a(i+1,l,n), \quad i \text{ odd}, \ i < l \le \lceil n/2 \rceil,$$
 (5.49)

$$a(i,l,n) < a(i+1,l,n), \quad i \text{ even}, \ i < l \le \lceil n/2 \rceil.$$
 (5.50)

PROOF. The proof is by induction on *i*. Separating the states based on activity of node 1 and node *n* yields the relations

$$a(i,l,n) = a(i-2,l-1,n-2) + a(i-1,l,n-1),$$
(5.51)

$$a(i,l,n) = a(i,l-1,n-2) + a(i,l,n-1),$$
(5.52)

,

with boundary conditions a(0, l, n) = 0 for all n and l, a(1, l, n) = 1 for l > 0 and all n and a(1, l, n) = 0 for $l \le 0$ and all n. Hence, the initialization step of the induction is

$$a(0,l,n) < a(1,l,n), \quad 0 < l < \lceil n/2 \rceil$$

 $a(0,l,n) = a(1,l,n), \quad l \le 0.$

Consider odd $i \le \lceil n/2 \rceil - 2$, let $i + 1 < l < \lceil n/2 \rceil$, and assume a(i, l, n) > a(i + 1, l, n). Using (5.51) and (5.52) we get

$$a(i+1,l,n) = a(i+1,l-1,n-2) + a(i+1,l,n-1)$$

< $a(i,l-1,n-2) + a(i+1,l,n-1) = a(i+2,l,n).$

This proves assertion (5.49). Assertions (5.48) and (5.50) can be proved in a similar manner. $\hfill \Box$

We now use Lemma 5.2 to prove Proposition 5.1.

PROOF. (Proposition 5.1) By relabeling the nodes in reverse order, we have that a(i, l, n) = a(n + i - 1, l, n). Using this, Assertion (i) can be shown by rewriting the throughput as follows:

$$\theta_i=Z_n^{-1}\sum_l a(i,l,n)\sigma^l=Z_n^{-1}\sum_l a(n-i+1,l,n)\sigma^l=\theta_{n-i+1}.$$

To prove assertion (ii) we first show that $(-1)^i(\theta_{i+1} - \theta_i)$ is positive. That is,

$$(-1)^{i}(\theta_{i+1} - \theta_{i}) = (-1)^{i} Z_{n}^{-1} \sum_{l}^{l} (a(i+1,l,n) - a(i,l,n)) \sigma^{l}$$
$$= 2(-1)^{i} Z_{n}^{-1} \sum_{l=i+1}^{\lfloor n/2 \rfloor} (a(i+1,l,n) - a(i,l,n)) \sigma^{l} > 0,$$
(5.53)

FAIRNESS IN LINEAR NETWORKS

where the inequality follows from Lemma 5.2. Proposition 5.1(ii) then follows from

$$\begin{aligned} (-1)^{i}(\theta_{i+1} - \theta_{i}) &= (-1)^{i} \Big(\theta_{i+1} - Z_{n}^{-1} \sum_{l} a(i,l,n) \sigma^{l} \Big) \\ &= (-1)^{i} \Big(\theta_{i+1} - Z_{n}^{-1} \sum_{l} (a(i,l-1,n-2) + a(i,l,n-1)) \sigma^{l} \Big) \\ &> (-1)^{i} \Big(\theta_{i+1} - Z_{n}^{-1} \sum_{l} (a(i,l-1,n-2) + a(i+1,l,n-1)) \sigma^{l} \Big) \\ &= (-1)^{i} \Big(\theta_{i+1} - Z_{n}^{-1} \sum_{l} a(i+2,l,n) \sigma^{l} \Big) = (-1)^{i+1} (\theta_{i+2} - \theta_{i+1}), \end{aligned}$$

where the second equality follows from (5.51), and the inequality from (5.53). $\hfill \Box$

6

ACHIEVING TARGET THROUGHPUTS

In this chapter we consider the CSMA model introduced in Section 1.3.2, and discuss the problem of determining the back-off rates that yield an arbitrary target throughput vector for general conflict graphs. To this end we study the throughput function that maps the back-off rates to the throughputs, and show that it is *globally invertible*. That is, every throughput vector inside the capacity region of the network can be achieved by a unique vector of back-off rates.

The present setting can be seen as a generalization of the problem considered in Chapter 5, where we focused on linear networks and equal throughputs. Explicit solutions for the inverse as obtained in Chapter 5 remain elusive for general conflict graphs and target throughputs. Instead we present three numerical methods to determine the inverse: fixed-point iteration, basic Newton iteration, and a continuation method (consisting of a sequence of Newton iteration steps).

This chapter is organized as follows. In Section 6.1 we introduce the model and describe the throughput function. Our main results on global invertibility are presented in Section 6.2. In Section 6.3 we describe several numerical methods for determining the inverse throughput function. Section 6.4 is concerned with results for special conflict graphs, including an alternative proof of Theorem 5.2 using Markov random fields. Finally, Section 6.5 presents some conclusions and a discussion.

ACHIEVING TARGET THROUGHPUTS

6.1 Model description

Consider the CSMA model from Section 1.3.2 on a general conflict graph. The backoff times and transmission durations are exponentially distributed. Since all results pertain to the stationary behavior of the CSMA model, we know by Theorem 4.1 that these remain valid for generally distributed back-off times and transmissions durations. The stationary distribution of the activity process (1.8) only depends on the ratio between transmission rates and back-off rates, so without loss of generality we can set $\mu_i \equiv 1$.

Denote the number of feasible states by K + 1, and write $\Omega = \{\omega_0, \omega_1, \dots, \omega_K\}$. The states are ordered such that $\omega_0 = \mathbf{0}$ (the empty state) and $\omega_k = \mathbf{e}_k$, the *k*th unit vector of \mathbb{R}^n , $k = 1, 2, \dots, n$. Note that the case K = n corresponds to the complete conflict graph, for which at most one node can be active at any time.

Recall that the stationary distribution of the activity process is denoted by π , with probability $\pi(\mathbf{y})$ as in (1.8). For the purpose of this chapter it is convenient to explicitly reflect the ordering of the states and the dependence on the back-off rate vector $\boldsymbol{\nu} = (\nu_1, \nu_2, \dots, \nu_n)^T$ in the notation, and introduce

$$\pi_k(\boldsymbol{\nu}) = \pi(\boldsymbol{\omega}_k) = \frac{\Lambda_k(\boldsymbol{\nu})}{Z(\boldsymbol{\nu})}, \quad k = 0, 1, \dots, K,$$

with $\boldsymbol{\omega}_k = (\boldsymbol{\omega}_{k,1}, \dots, \boldsymbol{\omega}_{k,n})^T$,

$$\Lambda_k(\nu) = \prod_{i=1}^n \nu_i^{\omega_{k,i}},\tag{6.1}$$

and $Z(\nu) = \sum_{k=0}^{K} \Lambda_k(\nu)$ the normalization constant.

We write the throughput of node *i* as $\theta_i(\nu)$, in order to explicitly reflect the dependence of the throughput on the back-off rates. The throughput vector $\theta(\nu) = (\theta_1(\nu), \theta_2(\nu), \dots, \theta_n(\nu))^T$ may be written as

$$\boldsymbol{\theta}(\boldsymbol{\nu}) = \sum_{k=0}^{K} \pi_k(\boldsymbol{\nu}) \boldsymbol{\omega}_k.$$

Recall from Section 1.3.5 that the range Γ of the mapping θ : $\mathbb{R}^n_+ \to \Gamma$ is the interior of the convex hull formed by all states $\omega_0, \omega_1, \ldots, \omega_K$, i.e.,

$$\Gamma = \operatorname{int} \Big\{ \sum_{k=0}^{K} \alpha_k \omega_k \Big| \sum_{k=0}^{K} \alpha_k = 1, \ \alpha_k \ge 0, k = 0, \dots, K \Big\}.$$

The problem of finding back-off rates that achieve a certain throughput vector can be formulated as finding $\nu_{\theta} = \nu_{\theta}(\gamma)$ that solves

$$\boldsymbol{\theta}(\boldsymbol{\nu}_{\boldsymbol{\theta}}) = \boldsymbol{\gamma},\tag{6.2}$$

with $\gamma = (\gamma_1, \dots, \gamma_n)^T \in \Gamma$. We thus need to study the mapping θ in detail.

6.2 GLOBAL INVERTIBILITY

6.2 Global invertibility

We first consider the non-normalized throughput

$$\eta(\boldsymbol{\nu}) = Z(\boldsymbol{\nu})\boldsymbol{\theta}(\boldsymbol{\nu}) = \sum_{k=0}^{K} \Lambda_k(\boldsymbol{\nu})\boldsymbol{\omega}_k.$$

This function is monotone in ν and hence easier to handle than the normalized throughput:

THEOREM 6.1. The mapping $\eta : \mathbb{R}^n_+ \to \mathbb{R}^n_+$ is globally invertible on its range \mathbb{R}^n_+ .

The proof of Theorem 6.1 is presented in Appendix 6.A.1. Theorem 6.1 says that the range of η is \mathbb{R}^n_+ , and that for any $\gamma \in \mathbb{R}^n_+$ we can find a unique $\nu_{\eta} = \nu_{\eta}(\gamma)$ that solves

$$\eta(\nu_{\eta}) = \gamma. \tag{6.3}$$

In some cases, it might be beneficial from a computational point of view to invert η rather than θ . Although η only represents the non-normalized throughput, this is sufficient when interested solely in the throughput ratios (for instance, when aiming for strict fairness).

The difference between η and θ is embodied by the normalization constant $Z(\nu)$, for which we have the following result.

LEMMA 6.1. Let $c \in \mathbb{R}^n_+$, s > 0 and write $\nu_{\eta}(sc)$ for the unique $\nu \in \mathbb{R}^n_+$ such that $\eta(\nu_{\eta}(sc)) = sc$. Then, the function $f_c(s) = s/Z(\nu_{\eta}(sc))$ is injective.

The proof of Lemma 6.1 is presented in Appendix 6.A.2. Lemma 6.1 suggests that we can control the throughput θ via the non-normalized throughput η , and indeed, it turns out to be a crucial ingredient in the proof of the following result.

THEOREM 6.2. The mapping $\theta : \mathbb{R}^n_+ \to \Gamma$ is globally invertible on Γ .

PROOF. It suffices to show that θ is injective. Let $\nu_1, \nu_2 \in \mathbb{R}^n_+$ be such that $\theta(\nu_1) = \theta(\nu_2)$. Then we have

$$\boldsymbol{\eta}(\boldsymbol{\nu}_1) = Z(\boldsymbol{\nu}_1)\boldsymbol{\theta}(\boldsymbol{\nu}_1), \quad \boldsymbol{\eta}(\boldsymbol{\nu}_2) = Z(\boldsymbol{\nu}_2)\boldsymbol{\theta}(\boldsymbol{\nu}_2). \tag{6.4}$$

With $c = \theta(\nu_1) = \theta(\nu_2) \in \mathbb{R}^n_+$, we consider the trajectory $\nu_{\eta}(sc)$, s > 0, for which we have

$$\eta(\boldsymbol{\nu}_{\boldsymbol{\eta}}(s\boldsymbol{c})) = s\boldsymbol{c} = Z(\boldsymbol{\nu}_{\boldsymbol{\eta}}(s\boldsymbol{c}))\boldsymbol{\theta}(\boldsymbol{\nu}_{\boldsymbol{\eta}}(s\boldsymbol{c})). \tag{6.5}$$

With $s_1 = Z(\nu_1)$, $s_2 = Z(\nu_2)$, it follows from Theorem 6.1 and (6.4), (6.5) that $\nu_{\eta}(s_1c) = \nu_1$, $\nu_{\eta}(s_2c) = \nu_2$, and that

$$\frac{1}{s_1}Z(\boldsymbol{\nu}_{\boldsymbol{\eta}}(s_1\boldsymbol{c})) = \frac{1}{s_2}Z(\boldsymbol{\nu}_{\boldsymbol{\eta}}(s_2\boldsymbol{c}))$$

Hence, by injectivity of $f_c(s)$ in Lemma 6.1, it follows that $s_1 = s_2$, so that $\nu_1 = \nu_2$.

Theorem 6.2 says that for any $\gamma \in \Gamma$, there is a unique vector $\nu_{\theta} = \nu_{\theta}(\gamma)$ that solves (6.2). The proofs of Theorems 6.1 and 6.2 require the description of the entire network structure, which appears at odds with the distributed nature of CSMA. However, in actual implementations, the back-off rates only have to be determined once, after which the nodes can operate fully autonomously. Thus, if the network structure is fixed, or if the time scale on which it changes is slower than that of the network operations, we retain a fully distributed CSMA protocol, while achieving the target throughputs.

6.3 Inversion methods

In Section 6.2 we established that both the non-normalized throughput η and the normalized throughput θ are globally invertible on their respective ranges. In this section we present several numerical procedures to compute the inverse of a given (normalized) throughput vector, as well as a light-traffic approximation of the throughput inverse.

6.3.1 Fixed-point iteration

A first numerical procedure to determine the inverse is fixed-point iteration. This procedure follows naturally from rewriting the system of non-linear equations (6.2) as a fixed-point equation. We distinguish between normalized throughput and non-normalized throughput.

Non-normalized throughput

Write

$$\theta_{i}(\boldsymbol{\nu}) = \sum_{k=0}^{n} \pi_{k}(\boldsymbol{\nu})\omega_{k,i} + \sum_{k=n+1}^{K} \pi_{k}(\boldsymbol{\nu})\omega_{k,i} = \nu_{i}\frac{1+G_{i}(\boldsymbol{\nu})}{Z(\boldsymbol{\nu})}, \quad (6.6)$$

with

$$G_i(\boldsymbol{\nu}) = \frac{1}{\nu_i} \sum_{k=n+1}^K \Lambda_k(\boldsymbol{\nu}) \omega_{k,i}.$$
(6.7)

We can thus write (6.3) as

$$\nu_{\eta} = \mathbf{H}(\nu_{\eta})$$

where

$$\mathbf{H}(\boldsymbol{\nu}) = \left(\frac{\gamma_i}{1 + G_i(\boldsymbol{\nu})}\right)_{i=1,\dots,n}$$

and G_i as in (6.7). Note that $\mathbf{H} : [\mathbf{0}, \gamma] \to [\mathbf{0}, \gamma]$, where we denote $[\mathbf{0}, \gamma] = [0, \gamma_1] \times \cdots \times [0, \gamma_n]$. By global invertibility of η , we know that ν_{η} is the unique fixed point that solves $\nu_{\eta} = \mathbf{H}(\nu_{\eta})$. Alternatively, since **H** is continuous, the existence of a fixed point also follows directly from Brouwer's fixed-point theorem.

The fixed-point iteration is defined as

$$\boldsymbol{\nu}_{\boldsymbol{\eta}}^{(0)} = \mathbf{0}, \quad \boldsymbol{\nu}_{\boldsymbol{\eta}}^{(l+1)} = \mathbf{H}(\boldsymbol{\nu}_{\boldsymbol{\eta}}^{(l)}), \quad l = 0, 1, \dots$$
 (6.8)

We next show that the iterands obtained through (6.8) approach the fixed point in a monotone fashion.

6.3 INVERSION METHODS

PROPOSITION 6.1. Assume that the conflict graph has no fully connected nodes (nodes that are connected to all the other nodes). Then, for i = 1, 2, ..., n and l = 1, 2, ..., n

$$0 = \nu_{\eta,i}^{(0)} < \nu_{\eta,i}^{(2)} < \dots < \nu_{\eta,i}^{(2l-2)} < \nu_{\eta,i} < \nu_{\eta,i}^{(2l-1)} < \nu_{\eta,i}^{(2l-3)} < \dots < \nu_{\eta,i}^{(3)} < \nu_{\eta,i}^{(1)} = \gamma_i.$$
(6.9)

PROOF. We have $\nu_{\eta}^{(0)} = \mathbf{0}$ by definition, $\nu_{\eta}^{(1)} = \gamma$ since $G_i(\mathbf{0}) = 0$, i = 1, ..., n, and $0 < \nu_{\eta,i} < \gamma_i$, i = 1, ..., n. Now let *l* be such that (6.9) holds for all i = 1, ..., n. Then $\nu_{\eta,i} < \nu_{\eta,i}^{(2l-1)}$, i = 1, ..., n, and by the exclusion of fully connected nodes we have that

$$G_i(\boldsymbol{\nu_\eta}) < G_i(\boldsymbol{\nu_\eta}^{(2l-1)}), \quad i=1,\ldots,n,$$

and so

$$H_i(\boldsymbol{\nu}_{\eta}) > H_i(\boldsymbol{\nu}_{\eta}^{(2l-1)}) = \boldsymbol{\nu}_{\eta,i}^{(2l)}, \quad i = 1, \dots, n$$

i.e.,

$$v_{n,i} > v_{n,i}^{(2l)}, \quad i = 1, \dots, n.$$
 (6.10)

In a similar fashion it follows from (6.10) that

$$v_{\eta,i} < v_{\eta,i}^{(2l+1)}, \quad i = 1, \dots, n.$$

The proof follows by induction.

Proposition 6.1 shows that the iteration scheme in (6.8) approaches the fixed point ever more closely, although it does not necessarily imply convergence.

Normalized throughput

We now present a similar fixed-point iteration scheme for $\nu_{\theta}(\gamma)$. Setting $\theta(\nu_{\theta}) = \gamma$ and rewriting (6.6) yields $\nu_{\theta} = \mathbf{K}(\nu_{\theta})$ with

$$\mathbf{K}(\boldsymbol{\nu}) = \left(\frac{\gamma_i Z(\boldsymbol{\nu})}{1 + G_i(\boldsymbol{\nu})}\right)_{i=1,\dots,n}.$$
(6.11)

We have thus established that $\nu_{\theta}(\gamma)$ is the unique solution to the fixed-point equation (6.11), and we can again try to find $\nu_{\theta}(\gamma)$ by iteration. That is, we let $\nu_{\theta}^{(0)} = \mathbf{0}$ and recursively define

$$\boldsymbol{\nu}_{\boldsymbol{\theta}}^{(l+1)} = \mathbf{K}(\boldsymbol{\nu}_{\boldsymbol{\theta}}^{(l)}), \quad l = 0, 1, \dots$$
(6.12)

To gain some insight into this fixed-point iteration, below we give two special cases for which we can prove convergence to the fixed point.

EXAMPLE 6.1. (Complete conflict graph) Assume that only one node may be active at any time. Let $\gamma = (\gamma, ..., \gamma)^T$, $\gamma \in \mathbb{R}_+$. By symmetry, both the solution $\nu_{\theta}(\gamma)$ as well as the iterands $\nu_{\theta}^{(l)}$, l = 0, 1, ... have identical components. Thus $Z(\nu_{\theta}) = 1 + n\nu_{\theta,1}$ and $G_i(\nu_{\theta}) = 0$. Iterating according to (6.12), gives for all i = 1, 2, ..., n,

$$v_{\theta,i}^{(l)} = \gamma(1+n\gamma+\cdots+(n\gamma)^{l-1}), \quad l=0,1,\ldots,$$

and $v_{\theta,i}(\gamma) = \lim_{l \to \infty} v_{\theta,i}^{(l)} = \frac{\gamma}{1-n\gamma}$ for $\gamma < \frac{1}{n}$. The requirement for convergence $\gamma < \frac{1}{n}$ is equivalent to $\gamma \in \Gamma$.

89

П

ACHIEVING TARGET THROUGHPUTS

In this particular example, $\nu_{\theta}(\gamma)$ can also be determined analytically. Noting that at most one node can be active at a time, and assuming all nodes to have the same back-off rate, it was shown in [92] that the throughput of node *i* equals $\theta_i = \nu_{\theta,1}/(1+n\nu_{\theta,1})$. Solving $\gamma = \theta_i$ for ν_{θ} then gives the same result as the fixed-point iteration.

EXAMPLE 6.2. **(Isolated nodes)** Assume that all nodes are isolated ($E = \emptyset$). As nodes do not interact, the throughput of node *i* equals $v_{\theta,i}/(1 + v_{\theta,i})$, and thus the choice $v_{\theta,i} = \gamma/(1 - \gamma)$, yields per-node throughputs $\gamma, \gamma < 1$.

The same result can be obtained by fixed-point iteration. Let $\gamma = (\gamma, ..., \gamma)^T$, $\gamma \in \mathbb{R}_+$, so the target vector, solution and iterands have identical components. We have

$$Z(\boldsymbol{\nu}_{\boldsymbol{\theta}}) = \binom{n}{0} + \binom{n}{1} \boldsymbol{\nu}_{\boldsymbol{\theta},1} + \binom{n}{2} \boldsymbol{\nu}_{\boldsymbol{\theta},1}^2 + \dots + \binom{n}{n} \boldsymbol{\nu}_{\boldsymbol{\theta},1}^n = (1 + \boldsymbol{\nu}_{\boldsymbol{\theta},1})^n$$

and

$$G_{i}(\boldsymbol{\nu}_{\boldsymbol{\theta}}) = \binom{n-1}{1} \boldsymbol{\nu}_{\theta,1} + \binom{n-1}{2} \boldsymbol{\nu}_{\theta,1}^{2} + \dots + \binom{n-1}{n-1} \boldsymbol{\nu}_{\theta,1}^{n-1}$$
$$= (1 + \boldsymbol{\nu}_{\theta,1})^{n-1} - 1,$$

so that

$$K_i(\boldsymbol{\nu}_{\boldsymbol{\theta}}) = \gamma(1 + \boldsymbol{\nu}_{\boldsymbol{\theta},1}). \tag{6.13}$$

By iterating (6.13), we get for all i = 1, 2, ..., n,

$$v_{\theta,i}^{(l)} = \gamma + \gamma^2 + \cdots + \gamma^l, \quad l = 0, 1, \ldots$$

Thus $v_{\theta,i}(\gamma) = \lim_{l \to \infty} v_{\theta,i}^{(l)} = \gamma/(1-\gamma)$, as expected.

Due to the inclusion of the normalization constant, the fixed-point iteration for the normalized throughput becomes theoretically less tractable than for the non-normalized throughput, and the counterpart of Proposition 6.1 remains elusive. In applying the iteration (6.12), though, we have encountered no convergence issues. See Section 6.3.3 for an example of a successful application of fixed-point iteration. In fact, for both the non-normalized and normalized throughputs the fixed-point iterations seems to work well.

6.3.2 Newton-based methods

A second numerical method for inverting the throughput function is Newton iteration. We present two versions: classical Newton iteration, and a continuation method. The latter method consists of a sequence of Newton iteration steps. Since there is no essential difference in these methods between the non-normalized and normalized case, we present the numerical procedures only for the normalized throughput θ .

Classical Newton iteration

Recall from basic Newton iteration that one selects an initial vector $\nu_{\theta}^{(0)} \in \Gamma$, and iterates according to

$$\boldsymbol{\nu}_{\boldsymbol{\theta}}^{(l+1)} = \boldsymbol{\nu}_{\boldsymbol{\theta}}^{(l)} - \left(\frac{\partial \boldsymbol{\theta}}{\partial \boldsymbol{\nu}}(\boldsymbol{\nu}_{\boldsymbol{\theta}}^{(l)})\right)^{-1} (\boldsymbol{\gamma} - \boldsymbol{\theta}(\boldsymbol{\nu}_{\boldsymbol{\theta}}^{(l)})), \quad l = 0, 1, \dots,$$

6.3 INVERSION METHODS

where

$$\frac{\partial \theta}{\partial \nu} = \left(\frac{\partial \theta_i}{\partial \nu_i}\right)_{i,j=1,\dots,n}$$

is the functional matrix, which also plays a crucial role in the proof of Theorem 6.1 (see (6.27) and further).

Continuation method

Let $\gamma = sc$, with $c \in \mathbb{R}^n_+$ and s > 0. The general idea behind the continuation method is to compute a sequence of back-off rates $\nu_{\theta}(\gamma^{(l)})$, with $\gamma^{(l)} = l\delta c$, and δ the step size such that s/δ is integer. Successive iterands are computed by performing a single Newton iteration step:

$$\boldsymbol{\nu}_{\boldsymbol{\theta}}^{(l+1)} = \boldsymbol{\nu}_{\boldsymbol{\theta}}^{(l)} - \left(\frac{\partial \boldsymbol{\theta}}{\partial \boldsymbol{\nu}}(\boldsymbol{\nu}_{\boldsymbol{\theta}}^{(l)})\right)^{-1} (l\delta \boldsymbol{c} - \boldsymbol{\theta}(\boldsymbol{\nu}_{\boldsymbol{\theta}}^{(l)})), \quad l = 0, \dots, s/\delta.$$

The step size affects the accuracy of the resulting approximation, as well as the computation time. In addition to finding $\nu_{\theta}(\gamma)$, the continuation method approximates the entire path $\nu_{\theta}(\gamma^{(l)})$, $l = 0, 1, ..., s/\delta$. Similar to the fixed-point iteration, both Newton-based methods can be modified to work for η as well. This is done in both cases by replacing the functional matrix $\frac{\partial \theta}{\partial \nu}$ by $\frac{\partial \eta}{\partial \nu}$ (see (6.19) and further).

6.3.3 Comparison of inversion methods

To illustrate the inversion methods we consider a linear network with n = 15 nodes in which an active node blocks all nodes within two hops. Note that for this particular network, we have a closed-form expression for the target back-off rates, see Theorem 5.2 and Proposition 6.3. We set as target throughput $\gamma = (1/5, ..., 1/5) \in \Gamma$. We apply each of the three inversions methods (fixed-point iteration, Newton iteration and continuation method) for 30 iterations, and compare in each step the back-off rates and throughputs to their respective target values. We measure the error of the iterands by the Euclidean norm. The results are shown in Figure 6.1. Figure 6.1(a) plots the error in the back-off rates, and Figure 6.1(b) shows the error in the corresponding throughputs. Both figures show convergence of all three methods.

In general, it is difficult to compare these methods, since the fixed-point method and the Newton-based methods have different computational bottlenecks. For the Newton-based methods, the initialization stage is the most cumbersome, in particular the computation of the matrix $\frac{\partial \theta}{\partial \nu}$ (or $\frac{\partial \eta}{\partial \nu}$). The iteration itself has a relatively low complexity. In contrast, the fixed-point method barely requires any initialization, but its iteration stage typically takes longer than that of the Newton-based methods (when aiming for equal accuracy). Thus, either method may be best, depending on the conflict graph, target throughput and required accuracy.

6.3.4 Light-traffic approximation

Starting from the fixed-point equation (6.11) we derive an approximation for the inverse $\nu_{\theta}(\gamma)$ which is accurate when the elements of the normalized throughput vector γ are relatively small (a similar result can be obtained for the non-normalized throughput).

ACHIEVING TARGET THROUGHPUTS



Figure 6.1: The relative errors of the target back-off rates and throughputs for fixedpoint iteration (black), Newton iteration (gray) and continuation method (dashed), plotted against the iteration step.

PROPOSITION 6.2. Let $\gamma \in \Gamma$ and denote the set of neighbors of node *i* by $\mathcal{N}_i = \{j : \{i, j\} \in E\}$. Then, as $\|\gamma\| \downarrow 0$,

$$v_{\theta,i}(\gamma) = \gamma_i(1+\gamma_i+\sum_{j\in\mathcal{N}_i}\gamma_j)+O(\|\gamma\|^3), \quad i=1,2,\ldots,n.$$

PROOF. For $\|\boldsymbol{\nu}\| \downarrow 0$, we know that $G_i(\boldsymbol{\nu}_{\boldsymbol{\theta}}) = O(\|\boldsymbol{\nu}_{\boldsymbol{\theta}}\|)$, and $Z(\boldsymbol{\nu}_{\boldsymbol{\theta}}) = 1 + O(\|\boldsymbol{\nu}_{\boldsymbol{\theta}}\|)$. Thus we obtain from (6.11)

$$v_{\theta,i} = \gamma_i + O(\|\boldsymbol{\gamma}\|^2),$$

and hence $\nu_{\theta} = \gamma + O(\|\gamma\|^2)$. Substituting this into (6.11) once more, and noting that $G_i(\gamma) = \sum_{j \neq i, j \notin \mathcal{N}_i} \gamma_j + O(\|\gamma\|^2)$, we deduce

$$\begin{split} \nu_{\theta,i} &= \frac{\gamma_i \big(1 + \sum_{j=1}^n \gamma_j + O\big(\|\gamma\|^2\big)\big)}{1 + G_i(\gamma) + O\big(\|\gamma\|^2\big)} = \gamma_i (1 + \sum_{j=1}^n \gamma_j) - \gamma_i G_i(\gamma) + O\big(\|\gamma\|^3\big) \\ &= \gamma_i (1 + \gamma_i + \sum_{j \in \mathcal{N}_i} \gamma_j) + O\big(\|\gamma\|^3\big), \quad \|\gamma\| \downarrow 0, \end{split}$$

as required.

The approximation in Proposition 6.2 may be understood by observing that a fraction of the time $\gamma_i + \sum_{j \in \mathcal{N}_i} \gamma_j + O(||\gamma||^2)$ node *i* is prevented from activating due to either its own activity or the activity of one of its neighbors.

6.4 Special conflict graphs

For certain specific conflict graphs, we can either find an exact expression for the fixed point $\nu_{\theta}(\gamma)$, or we can decompose the graph into several components in order to reduce the complexity of the inversion methods. We will exploit the fact that our model is a Markov random field. The crucial property of Markov random fields that we will use is that for any subset $S \subseteq V$, the distribution of S is determined by the

6.4 SPECIAL CONFLICT GRAPHS

state of its boundary, and is independent of all other nodes. That is, for a general conflict graph *G*, for any $\mathbf{y} \in \Omega$ and nonempty strict subset $S \subset V$, we have

$$\mathbb{P}(Y^{S} = \mathbf{y}^{S} | Y^{V \setminus S} = \mathbf{y}^{V \setminus S}) = \mathbb{P}(Y^{S} = \mathbf{y}^{S} | Y^{\partial S} = \mathbf{y}^{\partial S}),$$
(6.14)

with $Y^S = (Y_i)_{i \in S}$ the components of the vector *Y* indexed by *S*, and $\partial S = \{j \in V \setminus S : \{i, j\} \in E \text{ for some } i \in S\}$, the boundary of *S*.

The next proposition identifies the 'fair' back-off rates that render equal throughputs for all nodes in a linear topology. In particular, we consider *n* nodes on a line with a β -hop interference range, i.e., each node interferes with up to β adjacent nodes to the left and to the right, $n \ge \beta$. Note that this proposition is a reformulation of Theorem 5.2.

PROPOSITION 6.3. Consider the conflict graph that arises from the linear network described above, and let $\gamma = (\gamma, ..., \gamma)$ with $\gamma < 1/(\beta + 1)$. Then

$$\nu_{\theta,i}(\gamma) = \frac{\gamma(1-\beta\gamma)^{h_i-1}}{(1-(\beta+1)\gamma)^{h_i}},$$

with

$$h_{i} = \begin{cases} i, & i = 1, \dots, \beta, \\ \beta + 1, & i = \beta + 1, \dots, n - \beta, \\ n - i + 1, & i = n - \beta + 1, \dots, n, \end{cases}$$

the number of interferers of node i minus $\beta - 1$.

Note that $\gamma \to 1/(\beta + 1)$ as $v_{\theta,i}(\gamma) \to \infty$, so that the throughput approaches the maximum achievable fair throughput as the back-off rates tend to infinity. The proof of Proposition 6.3 can be found in Appendix 6.A.3. It is based on the Markov random field representation of the stationary distribution of the joint activity state, extending the approach in [39].

Before proceeding, we first introduce some additional notation. For any subset $S \subseteq V$, we may consider a modified version of the system with the nodes in $V \setminus S$ removed, or equivalently, a system associated with a conflict graph that is the subgraph of *G* induced by the nodes in *S* and the same back-off rates. For brevity, we will call such a modified version the *system induced by S*. Denote by Y(S) a random variable with the stationary distribution of the activity process in the system induced by *S* and by $\theta(S) = (\theta_v(S))_{v \in S}$ the associated throughput vector. Moreover, for any $S \subseteq V$, $W \subseteq V \setminus S$, let

$$\Delta(S;\mathbf{y}^W) = S \setminus \bigcup_{i \in W: y_i=1} \mathcal{N}_i$$

be the set of those nodes in *S* that are not blocked by nodes active under \mathbf{y}^W . By the definition of ∂S , we have that $\Delta(S; \mathbf{y}^{V \setminus S}) = \Delta(S; \mathbf{y}^{\partial S})$. Finally, let us denote by Ω^S the state space restricted to *S*.

Recall that $\mathcal{N}_i = \{j : \{i, j\} \in E\}$ is the set of neighbors of node *i* in the conflict graph *G*. We will now apply the property in (6.14) to show that the problem of finding the stationary distribution of *S* can be reduced to finding the stationary distribution of several smaller systems, by conditioning on the state of ∂S .

ACHIEVING TARGET THROUGHPUTS



Figure 6.2: A decomposable graph

PROPOSITION 6.4. For any conflict graph G = (V, E), $S \subseteq V$, and $\mathbf{y}^S \in \{0, 1\}^{|S|}$,

$$\mathbb{P}(Y^{S} = \mathbf{y}^{S}) = \sum_{\mathbf{y}^{\partial S} \in \Omega^{\partial S}} \mathbb{P}(Y(\Delta(S; \mathbf{y}^{\partial S})) = \mathbf{y}^{\Delta(S; \mathbf{y}^{\partial S})})$$
$$\cdot \mathbb{1}_{\{\sum_{i \in \partial S} \sum_{j \in \mathcal{N}_{i}} \mathcal{Y}_{i} \mathcal{Y}_{j} = 0\}} \mathbb{P}(Y^{\partial S} = \mathbf{y}^{\partial S}).$$
(6.15)

The proof of Proposition 6.4 is given in Appendix 6.A.4.

Proposition 6.4 may seem convoluted, but can be very useful in certain conflict graphs for reducing the complexity of solving inversion problems, by choosing the set *S* in a judicious way. For example, consider the conflict graph in Figure 6.2. In this case, the node set can be partitioned into two subsets V_1 and V_2 and a single node ν , so $V = V_1 \cup V_2 \cup \{\nu\}$. The sets V_1 and V_2 are not connected, and ν shares edges with nodes in both subgraphs. We can decompose the graphs V_1 and V_2 as follows.

COROLLARY 6.1. For any $\mathbf{y}^{V_1} \in \{0, 1\}^{|V_1|}$,

$$\mathbb{P}(Y^{V_1} = \mathbf{y}^{V_1}) = \mathbb{P}(Y(V_1) = \mathbf{y}^{V_1})(1 - \theta_{\mathcal{V}}(V)) + \mathbb{P}(Y(V_1 \setminus \mathcal{N}_{\mathcal{V}}) = \mathbf{y}^{V_1 \setminus \mathcal{N}_{\mathcal{V}}}) \mathbb{1}_{\{\mathbf{y}^{\mathcal{N}_{\mathcal{V}}} = 0\}} \theta_{\mathcal{V}}(V).$$

In particular, for any $i \in V_1 \cup \mathcal{N}_{v}$,

$$\theta_i(V) = \theta_i(V_1)(1 - \theta_v(V)), \tag{6.16}$$

and for any $i \in V_1 \setminus \mathcal{N}_{v}$,

$$\theta_i(V) = \theta_i(V_1)(1 - \theta_v(V)) + \theta_i(V_1 \setminus \mathcal{N}_v)\theta_v(V).$$
(6.17)

The proof of Corollary 6.1 is presented in Appendix 6.A.5.

If we now substitute $\theta_{\nu}(V) = \gamma_{\nu}$ into (6.16) and (6.17), then we see that the resulting inverse problem for finding $\nu_{\theta,i}$ only depends on the nodes in V_1 , and no longer requires knowledge about any node in V_2 . This allows us to solve the inversion problems for V_1 and V_2 separately. Doing so considerably reduces the complexity, as the number of feasible states of the induced subgraph on V_1 is much smaller than that of the entire graph. The result in Corollary 6.1 can also be applied when ν is replaced by a clique of nodes. Naturally, when the conflict graph is disconnected, each of the components can also be handled separately.

6.5 Concluding remarks

In this chapter we have established global invertibility of both the non-normalized and normalized throughput function for CSMA networks on general conflict graphs. This fundamental result, presented in Section 6.2, states that for any throughput vector inside the network capacity region there exists a unique vector of back-off rates that will lead to that throughput vector. This result allows us, for example, to compute the back-off rates that give equal throughputs among all nodes, or instead to create various user classes by designing the back-off rates so as to give certain nodes higher throughput than others. From Theorem 6.2 it immediately follows that the fair back-off rates obtained in Chapter 5 (for linear networks) are in fact unique.

In Section 6.3 we presented several algorithms for determining the back-off rates. The implementation of these algorithms involves the computation of the normalization constant $Z(\nu)$, the (inverse of) the functional matrix $\partial \theta / \partial \nu$, and the functions G_i in (6.7). These require the enumeration of the entire state space Ω , which essentially boils down to counting all independent sets of the conflict graph, a problem which is known to be computationally cumbersome for large graphs. An important task for future research is to find ways of dealing with this curse of dimensionality. One possible approach is to exploit the structure of the conflict graphs and using the theory of Markov random fields, as was done in Section 6.4. Another approach is to use the distributed algorithms in [34, 36].

Appendix

6.A Remaining proofs

6.A.1 Proof of Theorem 6.1

Rather than showing invertibility of η itself, we consider the mapping

$$\mathbf{f}(\mathbf{x}) = \ln \boldsymbol{\eta}(\mathbf{e}^{\mathbf{x}}), \quad \mathbf{x} = (x_1, x_2, \dots, x_n)^T \in \mathbb{R}^n,$$

with $e^{\mathbf{x}} = (e^{x_1}, e^{x_2}, \dots, e^{x_n})^T$ and $\ln \mathbf{y} = (\ln y_1, \dots, \ln y_n)^T$ for $\mathbf{y} = (y_1, y_2, \dots, y_n)^T \in \mathbb{R}^n$. Because ln and exp are invertible, global invertibility of **f** and η is equivalent.

By the main result in [99] we have that **f** is globally invertible if and only if (i) **f** is locally invertible and (ii) $\max_i |f_i(\mathbf{x})| \to \infty$ as $\max_i |x_i| \to \infty$.

To show that condition (i) holds, it suffices to show that the functional matrix

$$\frac{\partial \mathbf{f}}{\partial \mathbf{x}} = \left(\frac{\partial f_i}{\partial x_j}\right)_{i,j=1,\dots,n},$$

is non-singular at any point $\mathbf{x} \in \mathbb{R}^n$. Observe that

$$\frac{\partial f_i}{\partial x_j} = \frac{1}{\eta_i(\mathbf{e}^{\mathbf{x}})} \frac{\partial \eta_i(\mathbf{e}^{\mathbf{x}})}{\partial \mathbf{e}^{x_j}} \frac{\partial \mathbf{x}}{\partial x_j} = \frac{1}{\eta_i(\boldsymbol{\nu})} \frac{\partial \eta_i(\boldsymbol{\nu})}{\partial \nu_j} \nu_j.$$

$$\frac{\partial \mathbf{f}}{\partial \mathbf{x}} = \operatorname{diag}\left(\frac{1}{\eta_1(\nu)}, \dots, \frac{1}{\eta_n(\nu)}\right) \frac{\partial \eta}{\partial \nu} \operatorname{diag}(\nu_1, \dots, \nu_n), \tag{6.18}$$

Thus

with

96

$$\frac{\partial \eta}{\partial \nu} = \left(\frac{\partial \eta_i}{\partial \nu_j}\right)_{i,j=1,\dots,n}.$$

Because $v_1, \ldots, v_n > 0$, both diagonal matrices in (6.18) are non-singular, and we only have to verify that $\partial \eta / \partial \nu$ is non-singular as well.

By taking the derivative of $\Lambda_k(\nu)$, see (6.1), with respect to ν_j , we get

$$\frac{\partial \Lambda_k(\boldsymbol{\nu})}{\partial \nu_j} = \frac{1}{\nu_j} \Lambda_k(\boldsymbol{\nu}) \mathbb{1}_{\{\omega_{k,j}=1\}}, \quad k = 0, 1, \dots, K, \ j = 1, \dots, n.$$

Consequently,

$$\frac{\partial \eta_i}{\partial \nu_j} = \frac{1}{\nu_j} \sum_{k=0}^K \omega_{k,i} \omega_{k,j} \Lambda_k(\nu), \quad i, j = 1, 2, \dots, n.$$
(6.19)

Thus the functional matrix $\frac{\partial \eta}{\partial \nu}$ may be written as

$$\frac{\partial \boldsymbol{\eta}}{\partial \boldsymbol{\nu}} = \mathbf{P}(\boldsymbol{\nu})\mathbf{D}(\boldsymbol{\nu}),$$

with

$$\mathbf{P}(\boldsymbol{\nu}) = \sum_{k=0}^{K} \Lambda_k(\boldsymbol{\nu}) \boldsymbol{\omega}_k \boldsymbol{\omega}_k^T,$$

and

$$\mathbf{D}(\boldsymbol{\nu}) = \text{diag}(\nu_1^{-1}, \dots, \nu_n^{-1}). \tag{6.20}$$

The matrix **P** is positive definite since $\Lambda_k(\nu) > 0$, $\omega_k = \mathbf{e}_k$, k = 1, 2, ..., n. Therefore, $\frac{\partial \eta}{\partial \nu}$ is non-singular, as required.

In order to verify condition (ii), we write $\eta(\mathrm{e}^{\mathrm{x}})$ as

$$\eta_i(\mathbf{e}^{\mathbf{x}}) = \mathbf{e}^{x_i} \left(1 + \mathbf{e}^{-x_i} \sum_{k=n+1}^K \Lambda_k(\mathbf{e}^{\mathbf{x}}) \omega_{k,i} \right).$$
(6.21)

Let

$$m = \min_{i} e^{x_{i}}, \quad M = \max_{i} e^{x_{i}},$$
$$a = -\min_{i} x_{i}, \quad b = \max_{i} x_{i}.$$

It is seen from (6.1) and (6.21) that

$$\max_{i} \eta_i(\mathbf{e}^{\mathbf{x}}) \ge M = \mathbf{e}^b,\tag{6.22}$$

$$\min_{i} \eta_{i}(\mathbf{e}^{\mathbf{x}}) \le m(1 + (K - n)M^{n-1}) = \mathbf{e}^{-a}(1 + (K - n)\mathbf{e}^{(n-1)b}).$$
(6.23)

Assume that $\max_i |x_i| = \max\{a, b\} \to \infty$. We need to show that $\max_i |f_i(\mathbf{x})| \to \infty$ as well.

When $b \ge a$ we have

$$\max_{i} |f_{i}(\mathbf{e}^{\mathbf{x}})| \ge b = \max\{a, b\}.$$
(6.24)

6.A REMAINING PROOFS

When $b \le a$, we see from (6.22) and (6.23) that

$$\max_{i} |f_{i}(\mathbf{x})| \ge \max\{b, a - \ln(1 + (K - n)e^{(n-1)b})\} \ge \max\{b, a - A - Bb\},\$$

for some A, B > 0 only depending on K, n. Now

$$\min_{0 \le b^* \le a} \max\{b^*, a - A - Bb^*\} \ge b(a),$$

with b(a) the solution of b = a - A - Bb, i.e., $b(a) = \frac{a-A}{B+1}$. Hence, when $a \ge b$

$$\max_{i} |f_{i}(\mathbf{e}^{\mathbf{x}})| \ge \frac{a-A}{B+1} = \frac{\max\{a.b\} - A}{B+1}.$$
(6.25)

From (6.24) and (6.25) we see that

$$\max_{i} |f_i(\mathbf{e}^{\mathbf{x}})| \to \infty$$

as $\max\{a, b\} \to \infty$, and the proof is complete.

6.A.2 Proof of Lemma 6.1

In order to prove this lemma, we compute some derivatives. We have, see (6.19),

$$\frac{\partial Z(\boldsymbol{\nu})}{\partial v_j} = \sum_{k=0}^K \omega_{k,j} \frac{1}{v_j} \Lambda_k(\boldsymbol{\nu}) = \frac{1}{v_j} \eta_j(\boldsymbol{\nu}), \quad j = 1, \dots, n.$$

Recall from (6.19) that

$$\frac{\partial \eta_i}{\partial v_j} = \frac{1}{v_j} \sum_{k=0}^K \omega_{k,i} \omega_{k,j} \Lambda_k(\boldsymbol{\nu}), \quad i, j = 1, 2, \dots, n.$$

Differentiating $\eta(\nu_{\eta}(sc)) = sc$ with respect to *s*, we see that

$$\frac{\partial \eta}{\partial \nu}(\nu_{\eta}(sc))\cdot\nu'_{\eta}(sc)=c,$$

i.e.,

$$\boldsymbol{\nu}_{\boldsymbol{\eta}}'(s\boldsymbol{c}) = \left(\frac{\mathrm{d}(\boldsymbol{\nu}_{\boldsymbol{\eta}}(s\boldsymbol{c}))_1}{\mathrm{d}s}, \dots, \frac{\mathrm{d}(\boldsymbol{\nu}_{\boldsymbol{\eta}}(s\boldsymbol{c}))_n}{\mathrm{d}s}\right)^T = \left(\frac{\partial\boldsymbol{\eta}}{\partial\boldsymbol{\nu}}(\boldsymbol{\nu}_{\boldsymbol{\eta}}(s\boldsymbol{c}))\right)^{-1}\boldsymbol{c}.$$
 (6.26)

Moreover, we have

$$\frac{\partial \theta_i}{\partial \nu_j} = \frac{\partial}{\partial \nu_j} \left(\frac{\eta_i(\nu)}{Z(\nu)} \right) = \frac{1}{Z^2(\nu)} \left(\frac{\partial \eta_i(\nu)}{\partial \nu_j} Z(\nu) - \eta_i(\nu) \frac{\partial Z(\nu)}{\partial \nu_j} \right) \\ = \frac{1}{\nu_j} \left(\sum_{k=0}^K \pi_k(\nu) \omega_{k,i} \omega_{k,j} - \sum_{k=0}^K \pi_k(\nu) \omega_{k,i} \sum_{k=0}^K \pi_k(\nu) \omega_{k,j} \right).$$

Note that $\sum_{k=0}^{K} \pi_k(\nu) = 1$ and that $\theta(\nu) = \sum_{k=0}^{K} \pi_k(\nu)\omega_k$. Hence we have

$$\frac{\partial \theta_i}{\partial \nu_j} = \sum_{k=0}^K \pi_k(\nu) (\omega_k - \theta(\nu))_i (\omega_k - \theta(\nu))_j^T \frac{1}{\nu_j}.$$
(6.27)



ACHIEVING TARGET THROUGHPUTS

98

$$\frac{\partial \theta}{\partial \nu} = \mathbf{Q}(\nu) \mathbf{D}(\nu), \qquad (6.28)$$

with D the diagonal matrix in (6.20) and

$$\mathbf{Q}(\boldsymbol{\nu}) = \sum_{k=0}^{K} \pi_k(\boldsymbol{\nu}) (\boldsymbol{\omega}_k - \boldsymbol{\theta}(\boldsymbol{\nu})) (\boldsymbol{\omega}_k - \boldsymbol{\theta}(\boldsymbol{\nu}))^T.$$
(6.29)

From this characterization it is clear that $\mathbf{Q}(\nu)$ is positive semidefinite, and we will show below that this matrix is in fact positive definite. Assuming this, we compute from $\theta(\nu_{\eta}(sc)) = \frac{s}{Z(\nu_{\eta}(sc))}c$, for any s > 0

$$\frac{\partial \theta}{\partial \nu}(\nu_{\eta}(sc))\nu'_{\eta}(sc) = f'_{c}(s)c.$$
(6.30)

By (6.26) we have that $\nu'_{\eta}(sc) \neq 0$ and by the fact that $\mathbf{Q}(\nu)$ is positive definite and (6.28) we have that $\frac{\partial \theta}{\partial \nu}$ is non-singular at $\nu = \nu_{\eta}(sc)$. Hence, the left-hand side of (6.30) is a non-zero vector and so $f'_{c}(s)c \neq 0$. Hence $f'_{c}(s) \neq 0$ for any s > 0. Since $f_{c}(0) = 0$, $f_{c}(s) > 0$ for s > 0, the claim follows.

It remains to show that $\mathbf{Q}(\nu)$ is positive definite. Assume $\mathbf{y} \in \mathbb{R}^n$ is such that $\mathbf{Q}(\nu)\mathbf{y} = \mathbf{0}$. Then

$$0 = \mathbf{y}^T \mathbf{Q}(\boldsymbol{\nu}) \mathbf{y} = \sum_{k=0}^K \pi_k(\boldsymbol{\nu}) \left| (\boldsymbol{\omega}_k - \boldsymbol{\theta}(\boldsymbol{\nu}))^T \mathbf{y} \right|^2,$$

and so, as $\pi_k(\nu) > 0$, k = 0, 1, ..., K, we have

$$(\boldsymbol{\omega}_k - \boldsymbol{\theta}(\boldsymbol{\nu}))^T \mathbf{y} = 0, \quad k = 0, 1, \dots, K,$$

i.e.,

$$\boldsymbol{\omega}_{k}^{T}\mathbf{y} = \boldsymbol{\theta}(\boldsymbol{\nu})^{T}\mathbf{y}, \quad k = 0, 1, \dots, K.$$
(6.31)

Since $\omega_0 = 0$, we get $\theta(\nu)^T \mathbf{y} = 0$ from (6.31) with k = 0. Then, for k = 1, ..., n, it follows from $\omega_k = \mathbf{e}_k$ and (6.31) that

$$y_k = \boldsymbol{\omega}_k^T \mathbf{y} = \boldsymbol{\theta}(\boldsymbol{\nu})^T \mathbf{y} = 0.$$

Hence $\mathbf{y} = \mathbf{0}$. We conclude that $\mathbf{Q}(\boldsymbol{\nu})$ is non-singular, and then from (6.29) it is seen that $\mathbf{Q}(\boldsymbol{\nu})$ is positive definite.

6.A.3 Proof of Proposition 6.3

For conciseness, denote

$$\Psi_i = \mathbb{P}(Y_{i-\beta}, \dots, Y_{i-1} = 0), \quad i = \beta + 1, \dots, n+1,$$

and

$$a_i = \mathbb{P}(Y_i = 0 | Y_{i-\beta}, \dots, Y_{i-1} = 0), \quad i = \beta + 1, \dots, n$$

6.A REMAINING PROOFS

By definition,

$$\theta_{i} = \mathbb{P}(Y_{i} = 1) = \mathbb{P}(Y_{i} = 1, Y_{i-\beta}, \dots, Y_{i-1} = 0)$$

= $\mathbb{P}(Y_{i} = 1 | Y_{i-\beta}, \dots, Y_{i-1} = 0) \mathbb{P}(Y_{i-\beta}, \dots, Y_{i-1} = 0) = (1 - a_{i})\psi_{i},$ (6.32)

for all $i = \beta + 1, \ldots, n$.

The idea of the proof is to consider probabilities of the form $\mathbb{P}(Y_i = 1, Y_j = 0, j \in \mathcal{N}_i)$ and $\mathbb{P}(Y_i = 0, Y_j = 0, j \in \mathcal{N}_i)$ and use two different relationships between these in order to obtain a set of equations for the coefficients a_i .

First of all, it follows from the product form of the stationary distribution (or the local balance property) that

$$\mathbb{P}(Y_i = 1, Y_j = 0, j \in \mathcal{N}_i) = \nu_i \mathbb{P}(Y_i = 0, Y_j = 0, j \in \mathcal{N}_i),$$

for all $i = 1, \ldots, n$.

The second relationship between these two probabilities follows from the Markov random field representation of the stationary distribution.

Specifically, for all $i = 1, ..., \beta$, we may write

$$\mathbb{P}(Y_i = 0, Y_j = 0, j \in \mathcal{N}_i) = \mathbb{P}(Y_1, \dots, Y_\beta = 0) \prod_{l=\beta+1}^i \mathbb{P}(Y_i = 0 | Y_1, \dots, Y_{i-1} = 0)$$
$$= \psi_{\beta+1} \prod_{l=\beta+1}^i \mathbb{P}(Y_i = 0 | Y_{i-\beta}, \dots, Y_{i-1} = 0) = \psi_{\beta+1} \prod_{l=\beta+1}^i a_i.$$

For all $i = \beta + 1, \dots, n$, we may write

$$\mathbb{P}(Y_{i} = 0, Y_{j} = 0, j \in \mathcal{N}_{i})$$

$$= \mathbb{P}(Y_{i-\beta}, \dots, Y_{i-1} = 0) \prod_{l=i}^{\min\{i+\beta,n\}} \mathbb{P}(Y_{l} = 0 | Y_{i-\beta}, \dots, Y_{l-1} = 0)$$

$$= \psi_{i} \prod_{l=i}^{\min\{i+\beta,n\}} \mathbb{P}(Y_{l} = 0 | Y_{l-\beta}, \dots, Y_{l-1} = 0) = \psi_{i} \prod_{l=i}^{\min\{i+\beta,n\}} a_{i},$$

and

$$\begin{split} & \mathbb{P}(Y_i = 1, Y_j = 0, j \in \mathcal{N}_i) \\ &= \mathbb{P}(Y_{i-\beta}, \dots, Y_{i-1} = 0) \mathbb{P}(Y_i = 1 | Y_j = 0, j \in \mathcal{N}_i^-) \\ & \prod_{l=i+1}^{\min\{i+\beta,n\}} \mathbb{P}(Y_l = 0 | Y_i = 1, Y_{i-\beta}, \dots, Y_{i-1}, Y_{i+1}, \dots, Y_{l-1} = 0) = \psi_i (1 - a_i), \end{split}$$

yielding

$$\frac{\mathbb{P}(Y_i = 1, Y_j = 0, j \in \mathcal{N}_i)}{1 - a_i} = \frac{\mathbb{P}(Y_i = 0, Y_j = 0, j \in \mathcal{N}_i)}{a_i \prod_{j=i+1}^{\min\{i+\beta,n\}} a_j}$$

Now observe that $\psi_i + \sum_{j=i-\beta}^{i-1} \theta_j = 1$ for all $i = \beta + 1, ..., n$, and in particular $\psi_{\beta+1} + \sum_{j=1}^{\beta} \theta_j = 1$. Combining the above two sets of equations, we obtain

$$\theta_i = \nu_i (1 - \sum_{j=1}^{\beta} \theta_j) \prod_{j=\beta+1}^{i+\beta} a_j, \qquad (6.33)$$
ACHIEVING TARGET THROUGHPUTS

for $i = 1, \ldots, \beta$, while

$$1 - a_i = v_i a_i \prod_{j=i+1}^{\min\{i+\beta,n\}} a_j,$$
(6.34)

for all $i = \beta + 1, \ldots, n$.

A solution to (6.34) is provided by $a_i = a$ and $v_i = (1 - a)a^{-h_i}$, or equivalently $v_i = \alpha(1 + \alpha)^{h_i - 1}$, with $\alpha = 1/a - 1 > 0$. Taking $v_i = (1 - a)a^{-h_i}$, $i = 1, ..., \beta$, we obtain from (6.33)

$$\theta_1 = \cdots = \theta_\beta = \theta = (1-a)(1-\beta\theta),$$

i.e.,

$$\theta_1 = \cdots = \theta_\beta = \theta = \frac{1-a}{1+\beta(1-a)} = \frac{\alpha}{1+(\beta+1)\alpha},$$

and (6.32) then yields

$$\theta = \frac{1-a}{1+\beta(1-a)} = \frac{\alpha}{1+(\beta+1)\alpha} \quad \text{for all} \quad i = \beta+1, \dots, n.$$

Noting that

$$a = \frac{1 - (\beta + 1)\theta}{1 - \beta\theta}$$
 or $\alpha = \frac{\theta}{1 - (\beta + 1)\theta}$

then completes the proof.

6.A.4 Proof of Proposition 6.4

The product form of the stationary distribution implies

$$\mathbb{P}(Y^{S} = \mathbf{y}^{S} \mid Y^{V \setminus S} = \mathbf{y}^{V \setminus S})$$

$$= \frac{\mathbb{P}(Y^{S} = \mathbf{y}^{S}, Y^{V \setminus S} = \mathbf{y}^{V \setminus S})}{\mathbb{P}(Y^{V \setminus S} = \mathbf{y}^{V \setminus S})} = \frac{\mathbb{P}(Y = \mathbf{y})}{\sum_{(\mathbf{x}^{S}, \mathbf{y}^{V \setminus S}) \in \Omega} \mathbb{P}(Y^{S} = \mathbf{x}^{S}, Y^{V \setminus S} = \mathbf{y}^{V \setminus S})}$$

$$= \frac{Z^{-1} \prod_{j=1}^{n} v_{j}^{y_{j}}}{Z^{-1} \sum_{(\mathbf{x}^{S}, \mathbf{y}^{V \setminus S}) \in \Omega} \prod_{i \in S} v_{i}^{x_{i}} \prod_{i \notin S} v_{j}^{y_{j}}} = K^{-1}(S; \mathbf{y}^{V \setminus S}) \prod_{j \in S} v_{j}^{y_{j}},$$

for any $\mathbf{y} \in \Omega$, with

$$K(S; \mathbf{y}^{V \setminus S}) = \sum_{\mathbf{x}^{S}: (\mathbf{x}^{S}, \mathbf{y}^{V \setminus S}) \in \Omega} \prod_{i \in S} \nu_{i}^{x_{i}}.$$

Note that

$$\prod_{j\in S} \nu_j^{\gamma_j} = \prod_{j\in \Delta(S;\mathbf{y}^{\partial S})} \nu_j^{\gamma_j}$$

for any $\mathbf{y} \in \Omega$. Likewise,

$$\prod_{i \in S} v_i^{x_i} = \prod_{i \in \Delta(S; \mathbf{y}^{\partial S})} v_i^{x_i}$$

for any $(\mathbf{x}^{S}, \mathbf{y}^{V \setminus S}) \in \Omega$.

6.A REMAINING PROOFS

Let $\Omega(S; \mathbf{y}^{V \setminus S})$ be the collection of independent sets in the subgraph of *G* induced by $\Delta(S; \mathbf{y}^{\partial S})$. It is easily verified that $(\mathbf{x}^S, \mathbf{y}^{V \setminus S}) \in \Omega$ if and only if it holds that $\mathbf{x}^{\Delta(S; \mathbf{y}^{\partial S})} \in \Omega(S; \mathbf{y}^{\partial S})$. It follows that

$$K(S; \mathbf{y}^{V \setminus S}) = \sum_{\mathbf{x}^{\Delta(S; \mathbf{y}^{\partial S})} \in \Omega(S; \mathbf{y}^{\partial S})} \prod_{i \in \Delta(S; \mathbf{y}^{\partial S})} \nu_i^{x_i},$$

and thus corresponds to the normalization constant of the system induced by $\Delta(S; \mathbf{y}^{\partial S})$.

...

We conclude that

$$\mathbb{P}(Y^{S} = \mathbf{y}^{S} | Y^{V \setminus S} = \mathbf{y}^{V \setminus S}) = \frac{\prod_{j \in \Delta(S; \mathbf{y}^{\partial S})} v_{j}^{Y_{j}}}{\sum_{\mathbf{x}^{\Delta(S; \mathbf{y}^{\partial S})} \in \Omega(S; \mathbf{y}^{\partial S})} \prod_{i \in \Delta(S; \mathbf{y}^{\partial S})} v_{i}^{X_{i}}}$$
$$= \mathbb{P}(Y(\Delta(S; \mathbf{y}^{\partial S})) = \mathbf{y}^{\Delta(S; \mathbf{y}^{\partial S})}),$$

for any $\mathbf{y} \in \Omega$. Informally speaking, the distribution of the activity state of the nodes in *S* in the original system, conditional on the activity states of the remaining nodes, equals the stationary distribution of the system induced by $\Delta(S; \mathbf{y}^{\partial S})$. Since $\Delta(S; \mathbf{y}^{\partial S})$ only depends on $\mathbf{y}^{V\setminus S}$ through $\mathbf{y}^{\partial S}$, it further follows that $\mathbb{P}(Y^S = \mathbf{y}^S | Y^{V\setminus S} = \mathbf{y}^{V\setminus S}) =$ $\mathbb{P}(Y^S = \mathbf{y}^S | Y^{\partial S} = \mathbf{y}^{\partial S})$. This corroborates the fact that the stationary distribution is a Markov random field with a neighborhood structure defined by the conflict graph *G*.

Now observe that

$$\mathbb{P}(Y^{S} = \mathbf{y}^{S} | Y^{V \setminus S} = \mathbf{y}^{V \setminus S}) = \mathbb{P}(Y^{S} = \mathbf{y}^{S} | Y^{\partial S} = \mathbf{y}^{\partial S}) = 0$$

unless

$$\sum_{i\in\partial S}\sum_{j\in\mathcal{N}_i}y_iy_j=0.$$

Thus we may write

$$\mathbb{P}(Y^{S} = \mathbf{y}^{S} | Y^{\partial S} = \mathbf{y}^{\partial S}) = \mathbb{P}(Y(\Delta(S; \mathbf{y}^{\partial S})) = \mathbf{y}^{\Delta(S; \mathbf{y}^{\partial S})}) \mathbb{1}_{\{\sum_{i \in \partial S} \sum_{j \in \mathcal{N}_{i}} y_{i} y_{j} = 0\}},$$

for all $\mathbf{y} \in \{0, 1\}^V$, rather than just $\mathbf{y} \in \Omega$.

We deduce that

$$\mathbb{P}(Y^{S} = \mathbf{y}^{S}) = \sum_{\mathbf{y}^{\partial S} \in \Omega^{\partial S}} \mathbb{P}(Y^{S} = \mathbf{y}^{S} | Y^{\partial S} = \mathbf{y}^{\partial S}) \mathbb{P}(Y^{\partial S} = \mathbf{y}^{\partial S})$$
$$= \sum_{\mathbf{y}^{\partial S} \in \Omega^{\partial S}} \mathbb{P}(Y(\Delta(S; \mathbf{y}^{\partial S})) = \mathbf{y}^{\Delta(S; \mathbf{y}^{\partial S})}) \mathbb{1}_{\{\sum_{i \in \partial S} \sum_{j \in \mathcal{N}_{i}} y_{i} y_{j} = 0\}} \mathbb{P}(Y^{\partial S} = \mathbf{y}^{\partial S}).$$
(6.35)

6.A.5 Proof of Corollary 6.1

For the specific graph under consideration, apply Proposition 6.4 with $S = V_1$ so that $\partial S = \{v\}$. We have

$$\begin{split} \mathbb{P}(Y^{V_1} = \mathbf{y}^{V_1} | Y^{V \setminus V_1} = \mathbf{y}^{V \setminus V_1}) &= \mathbb{P}(Y^{V_1} = \mathbf{y}^{V_1} | Y_{\mathcal{V}} = y_{\mathcal{V}}) \mathbb{1}_{\{y_{\mathcal{V}} \sum_{j \in \mathcal{N}_{\mathcal{V}}} y_j = 0\}} \\ &= \begin{cases} \mathbb{P}(Y(V_1) = \mathbf{y}^{V_1}), & y_{\mathcal{V}} = 0, \\ \mathbb{P}(Y(V_1 \setminus \mathcal{N}_{\mathcal{V}}) = \mathbf{y}^{V_1 \setminus \mathcal{N}_{\mathcal{V}}}) \mathbb{1}_{\{\mathbf{y}^{\mathcal{N}_{\mathcal{V}}} = 0\}}, & y_{\mathcal{V}} = 1, \end{cases} \end{split}$$

102 ACHIEVING TARGET THROUGHPUTS

with $\mathbb{P}(Y_{\nu} = 1) = 1 - \mathbb{P}(Y_{\nu} = 0) = \theta_{\nu}(V)$, so that (6.35) reduces to

$$\mathbb{P}(Y^{V_1} = \mathbf{y}^{V_1}) = \mathbb{P}(Y(V_1) = \mathbf{y}^{V_1})(1 - \theta_{\mathcal{V}}(V)) + \mathbb{P}(Y(V_1 \setminus \mathcal{N}_{\mathcal{V}}) = \mathbf{y}^{V_1 \setminus \mathcal{N}_{\mathcal{V}}})\mathbb{1}_{\{\mathbf{y}^{\mathcal{N}_{\mathcal{V}}} = 0\}}\theta_{\mathcal{V}}(V).$$

In particular, we obtain

$$\begin{aligned} \theta_{i}(V) &= \sum_{\mathbf{y}^{V_{1}}: y_{i}=1} \mathbb{P}(Y^{V_{1}} = \mathbf{y}^{V_{1}}) \\ &= \sum_{\mathbf{y}^{V_{1}}: y_{i}=1} \left[\mathbb{P}(Y(V_{1}) = \mathbf{y}^{V_{1}}) (1 - \theta_{v}(V)) \\ &+ \mathbb{P}(Y(V_{1} \setminus \mathcal{N}_{v}) = \mathbf{y}^{V_{1} \setminus \mathcal{N}_{v}}) \mathbb{1}_{\{\mathbf{y}^{\mathcal{N}_{v}} = 0\}} \theta_{v}(V) \right] \\ &= \begin{cases} \theta_{i}(V_{1}) (1 - \theta_{v}(V)), & i \in \mathcal{N}_{v}, \\ \theta_{i}(V_{1}) (1 - \theta_{v}(V)) + \theta_{i}(V_{1} \setminus \mathcal{N}_{v}) \theta_{v}(V), & i \notin \mathcal{N}_{v}. \end{cases} \end{aligned}$$

7

OPTIMAL TRADEOFF BETWEEN EXPOSED AND HIDDEN NODES

In this chapter we adapt the CSMA model introduced in Section 1.3.2 to incorporate collisions, and we evaluate the impact of the carrier-sensing range on the network performance. The effect of the sensing range can be understood as follows. A small range allows for more simultaneous transmissions, but is less effective in reducing collisions (hidden nodes). On the other hand, a large sensing range mitigates interference, but also admits fewer transmissions (exposed nodes). The model considered in the present chapter differs slightly from that in Chapters 4-6, where it was assumed that the carrier-sensing mechanism precludes all collisions. In contrast to Chapters 4-6, we assume that the back-off rates and transmission rates of all nodes are fixed and equal.

The main contribution of this chapter is the examination of the impact of hidden and exposed nodes on the throughput. We consider a linear network in the asymptotic regime where the number of nodes in the network tends to infinity. For such networks we are able to obtain structural results on the joint effect of hidden nodes and exposed nodes. We determine analytically the throughput-optimal sensing range that achieves the best tradeoff between reducing hidden nodes and preventing exposed nodes.

This chapter is structured as follows. In Section 7.1 we introduce the model, and derive some auxiliary results. Section 7.2 discusses the main results on the carriersensing tradeoff. In Section 7.3 we perform a detailed study of the normalization constant. In Section 7.4 we validate the analytical results for the linear network by simulation, and investigate networks with more general topologies, and in Section 7.5 we present some concluding remarks. OPTIMAL TRADEOFF BETWEEN EXPOSED AND HIDDEN NODES

7.1 Model description and preliminary results

We again consider the CSMA model introduced in Section 1.3.2. The network consists of a linear array of 2n + 1 evenly spaced nodes with sensing range β , and we denote the set of all nodes by $\mathcal{N} = \{-n, ..., n\}$. Since we aim to model collisions in this chapter, we have to take the destination of each transmission into account. Whenever a node activates, it transmits a single packet to a neighboring node. With probability ψ , the packet is intended for its right neighbor, and with probability $1 - \psi$ for its left neighbor. To accommodate this, we introduce (pure destination) nodes n + 1 and -(n + 1), which receive packets, but do not transmit packets themselves. As will be shown in Proposition 7.2, the throughput is insensitive to the parameter ψ .

The length of the back-off period is assumed to be exponentially distributed with mean $1/\sigma$, while transmissions last for an exponentially distributed duration with unit mean. As all results in this chapter are based on the stationary behavior of the activity process, we know by Theorem 4.1 that these results in fact hold for generally distributed back-off times and transmission durations.

Each state of the activity process is described by

$$\omega = (\omega_{-n}, \dots, \omega_n) \in \{0, 1\}^{2n+1}.$$

This process has stationary distribution (see (1.8))

$$\pi(\omega) = \begin{cases} Z_{2n+1}^{-1} \prod_{\nu=-n}^{n} \sigma^{\omega_{\nu}}, & \text{if } \omega \text{ is feasible,} \\ 0, & \text{otherwise,} \end{cases}$$
(7.1)

with Z_{2n+1} the normalization constant. The normalization constant can be defined recursively as ([11, 71])

$$Z_i = \begin{cases} 1+i\sigma, & i=0,1,\ldots,\beta+1\\ Z_{i-1}+\sigma Z_{i-\beta-1}, & i\geq \beta+2. \end{cases}$$

The sequence $(Z_i)_{i=0}^{\infty}$ is well studied. In fact, for a network with *i* nodes, Z_i represents the normalization constant. Straightforward calculations show that the generating function $G_Z(x)$ of Z_i can be written as (see [71])

$$G_Z(x) = \sum_{i=0}^{\infty} Z_i x^i = \frac{x - 1 + \sigma x^{\beta + 1} - \sigma x}{(x - 1)(1 - x - \sigma x^{\beta + 1})}.$$
(7.2)

Let $\lambda_0, \ldots, \lambda_\beta$ denote the $\beta + 1$ distinct roots (see Proposition 7.8) of

$$\lambda^{\beta+1} - \lambda^{\beta} - \sigma = 0. \tag{7.3}$$

We denote by λ_0 the unique positive real root for which $\lambda_0 > |\lambda_j|$, $j \neq 0$ (see [71]). Applying partial fraction expansion to (7.2) yields the following result:

PROPOSITION 7.1. The normalization constant Z_i is given by

$$Z_{i} = \sum_{j=0}^{\beta} c_{j} \lambda_{j}^{i} \quad , i = 0, 1, \dots,$$
(7.4)

7.1 MODEL DESCRIPTION AND PRELIMINARY RESULTS

where λ_j are the roots of (7.3), and

$$c_j = \frac{\lambda_j^{\beta+1}}{(\beta+1)\lambda_j - \beta}$$

The proof of Proposition 7.1 is provided in Appendix 7.A, along with the other proofs not given in the main text. The representation (7.4) was also used in Chapter 5. Note that Proposition 7.1 does not rely on previous results.

We use the protocol model discussed in Section 1.1.1 to describe interference. Since all nodes are evenly spaced, this model gives rise to an interference range η . We assume that a transmission succeeds if and only if at the start of this transmission no nodes within distance η of the receiving node are already active. This type of interference is referred to in the literature as the *perfect capture* collision model [11]. Note that (7.1) does not depend on η , as collisions have no impact on the dynamics of the system. Using the sensing range β and interference range η we can formally define hidden nodes and exposed nodes. Consider a transmission from node v to node w. Hidden nodes are then defined as nodes that are outside the sensing range of v, but within the interference range of w. Such nodes are not blocked by the activity of node v, but their proximity to node w makes the hidden nodes harmful to the transmission from v to w. Conversely, exposed nodes are those nodes that are within the sensing range of v, but outside the interference range of w. Such nodes are blocked by an ongoing transmission from ν to w, despite the fact that they will not cause this transmission to fail. Denote by \mathcal{H}_r (\mathcal{H}_l) the set of hidden nodes of transmissions from node 0 to node 1 (node -1): all nodes outside the sensing range of 0, but within the interference range of the receiving node 1 (node -1). By \mathcal{E}_r (\mathcal{E}_l) we denote the set of nodes to which this transmission is exposed, so all nodes within the sensing range of 0, but outside the interference range of the receiving node. For completeness we let \mathcal{B}_r (\mathcal{B}_l) denote the set of all remaining nodes that block transmissions from node 0 to node 1 (node -1). This yields:

$$\begin{aligned} \mathcal{H}_r &= \{ \nu \in \mathcal{N} \mid |\nu| \geq \beta + 1, \ |\nu - 1| \leq \eta \}, \quad \mathcal{H}_l = \{ \nu \in \mathcal{N} \mid |\nu| \geq \beta + 1, \ |\nu + 1| \leq \eta \}, \\ \mathcal{E}_r &= \{ \nu \in \mathcal{N} \mid |\nu| \leq \beta, \ |\nu - 1| \geq \eta + 1 \}, \quad \mathcal{E}_l = \{ \nu \in \mathcal{N} \mid |\nu| \leq \beta, \ |\nu + 1| \geq \eta + 1 \}, \\ \mathcal{B}_r &= \{ \nu \in \mathcal{N} \mid |\nu| \leq \beta, \ |\nu - 1| \leq \eta \}, \quad \mathcal{B}_l = \{ \nu \in \mathcal{N} \mid |\nu| \leq \beta, \ |\nu + 1| \leq \eta \}. \end{aligned}$$

So $\mathcal{E}_r \cup \mathcal{B}_r = \mathcal{E}_l \cup \mathcal{B}_l = \{v \in \mathcal{N} \mid |v| \le \beta\}$. An example is given in Figure 7.1(a). Node 3 is a hidden node, as it interferes with the transmission from node 0 to node 1 ($\eta = 2$) despite the carrier-sensing mechanism ($\beta = 1$). In Figure 7.1(b) node 0 is an exposed node to the transmission from node 2 to node 3 because it would not interfere ($\eta = 2$) with this transmission but is nevertheless silenced by the activity of node 2 ($\beta = 2$).

We focus on node 0 (the node in the middle of the network) and in particular its throughput $\theta_n(\beta, \eta, \sigma)$ defined as the average number of successful transmissions per unit of time.

PROPOSITION 7.2. The throughput of node 0 is given by

$$\theta_n(\beta,\eta,\sigma) = \sigma \frac{Z_{n-\max\{\beta,\eta-1\}} Z_{n-\max\{\beta,\eta+1\}}}{Z_{2n+1}}.$$
(7.5)

OPTIMAL TRADEOFF BETWEEN EXPOSED AND HIDDEN NODES



(a) Node 3 is a hidden node, and may interfere with the transmission between nodes 0 and 1.



β

Figure 7.1: Examples of hidden and exposed nodes.

PROOF. Denote by θ_r (θ_l) the rate of successful transmission of node 0 to node 1 (node -1), so $\theta_n(\beta, \eta, \sigma) = \theta_r + \theta_l$. The activation attempts to node 1 (node -1) occur according to a Poisson process with rate $\sigma \psi$ (rate $\sigma(1 - \psi)$). We first consider activation attempts to node 1. Whether an activation attempt is successful depends on the state of the system when this attempt occurs. Define

$$A_{1} = \{ \omega \in \Omega \mid \exists v \in \mathcal{B}_{r} \cup \mathcal{E}_{r} : \omega_{v} = 1 \},\$$

$$A_{2} = \{ \omega \in \Omega \mid \forall v \in \mathcal{B}_{r} \cup \mathcal{E}_{r} : \omega_{v} = 0, \exists v \in \mathcal{H}_{r} : \omega_{v} = 1 \},\$$

$$A_{3} = \{ \omega \in \Omega \mid \forall v \in \mathcal{B}_{r} \cup \mathcal{E}_{r} \cup \mathcal{H}_{r} : \omega_{v} = 0 \}.$$

When the system is in state $\omega \in A_1$, the attempt is blocked and node 0 remains in its current state. When the system is in a state $\omega \in A_2$, node 0 is not blocked so it activates. However, at least one hidden node is active so the transmission fails and does not contribute to the throughput. When the system is in state $\omega \in A_3$, the perfect capture assumption guarantees a successful transmission. It follows from the PASTA property (cf. [5]) that the probability of an arbitrary activation attempt resulting in a successful transmission is equal to the limiting probability of the system being in a state $\omega \in A_3$. So the rate of successful transmissions initiated (and thus the throughput) is given by

$$\theta_r = \sigma \psi \sum_{\omega \in A_3} \pi(\omega).$$

From the definitions of \mathcal{B}_r , \mathcal{E}_r and \mathcal{H}_r we see that

$$A_3 = \{ \omega \in \Omega \mid \forall v \in (D_1 \cup D_2)^c : \omega_v = 0 \},\$$

where

 $D_1 = \{-n, \ldots, -\max\{\beta, \eta - 1\} - 1\}, \quad D_2 = \{\max\{\beta, \eta + 1\} + 1, \ldots, n\}.$

Let Z_D denote the normalization constant for a subset of nodes $D \subseteq \mathcal{N}$ defined as

$$Z_D = \sum_{\omega \in \Omega, \ \forall \nu \in D^c: \omega_{\nu} = 0} \prod_{\nu = -n}^n \sigma^{\omega_{\nu}}.$$

Then

$$\theta_r = \sigma \psi \frac{Z_{D_1 \cup D_2}}{Z_N}$$

The model on the line has the property that by conditioning on the activity of one of the nodes, its state space can be decomposed, leading to two smaller instances of the same model on the line. In particular, we know that $Z_{D_1 \cup D_2} = Z_{D_1} Z_{D_2}$ (see [11, Equation (15)]), so that

$$\theta_r = \sigma \psi \frac{Z_{D_1} Z_{D_2}}{Z_{\mathcal{N}}} = \sigma \psi \frac{Z_{n-\max\{\beta,\eta-1\}} Z_{n-\max\{\beta,\eta+1\}}}{Z_{2n+1}},$$

where Z_i denotes the normalization constant of a network with *i* consecutive nodes on a line. Similarly,

$$\theta_l = \sigma(1-\psi) \frac{Z_{n-\max\{\beta,\eta-1\}} Z_{n-\max\{\beta,\eta+1\}}}{Z_{2n+1}}.$$

and (7.5) follows by adding θ_r and θ_l .

7.2 Main results

Our principal aim is to choose the sensing range β so that the throughput $\theta_n(\beta, \eta, \sigma)$ is maximized for a given η and σ . Define

$$\beta_n^* = \arg \max_{\beta} \theta_n(\beta, \eta, \sigma).$$

Determining β_n^* corresponds to quantifying and optimizing the tradeoff between preventing collisions through interference (preventing hidden nodes by setting β large) and allowing harmless transmissions (preventing exposed nodes by setting β small). We want to obtain structural insights in how to choose β_n^* , and for this purpose the expressions for Z_i in (7.4) and $\theta_n(\beta, \eta, \sigma)$ in (7.5) are too cumbersome. Therefore, we investigate the throughput in the regime where the network becomes large $(n \to \infty)$, so that (7.5) simplifies considerably.

The analytic results that we obtain for the infinite network provide remarkably sharp approximations for the finite network; see Section 7.4.1.

We start by presenting the limiting expression for $\theta_n(\beta, \eta, \sigma)$ as the size of the network grows large:

PROPOSITION 7.3. Let λ_0 denote the unique positive real root of (7.3). Then

$$\theta(\beta,\eta,\sigma) = \lim_{n \to \infty} \theta_n(\beta,\eta,\sigma) = \sigma \frac{\lambda_0^{\beta-f(\beta)}}{(\beta+1)\lambda_0 - \beta},$$
(7.6)

where

$$f(\beta) = \begin{cases} 2\eta, & \text{if } 0 \le \beta \le \eta - 1, \\ \eta + \beta + 1, & \text{if } \eta - 1 \le \beta \le \eta + 1, \\ 2\beta, & \text{if } \beta \ge \eta + 1. \end{cases}$$
(7.7)

Optimal tradeoff between exposed and hidden nodes

PROOF. From Rouché's theorem (see [16]) it readily follows that $\lambda_0 > |\lambda_j|$ for $j = 1, ..., \beta$, and so from (7.4) we get

$$Z_i = c_0 \lambda_0^i \left(1 + o(1) \right), \quad i \to \infty$$

Hence

$$\lim_{n \to \infty} \theta_n(\beta, \eta, \sigma) = \lim_{n \to \infty} \sigma \frac{c_0 \lambda_0^{n - \max\{\beta, \eta - 1\}} c_0 \lambda_0^{n - \max\{\beta, \eta + 1\}}}{c_0 \lambda_0^{2n + 1}}$$
$$= \sigma c_0 \lambda_0^{- \max\{\beta, \eta - 1\} - \max\{\beta, \eta + 1\} - 1},$$

which yields (7.7).

Now that we have the limiting expression for the throughput in (7.6), we opt for an asymptotic analysis. That is, instead of searching for β_n^* , we search for its asymptotic counterpart

$$\beta^* = \arg\max_{\beta} \theta(\beta, \eta, \sigma),$$

where we henceforth consider θ as a function of the real variable $\beta \ge 0$. In Section 7.4.1 we show that the errors $|\theta_n - \theta|$ and $|\beta_n^* - \beta^*|$ become small, already for moderate values of *n*. Because we consider from here onwards the regime $n \to \infty$, all nodes have the same number of nodes within their sensing range. This removes all boundary effects, and all nodes have the same throughput, which is why just investigating node 0 is sufficient to investigate the entire network.

PROPOSITION 7.4. $\beta^* \in [\eta - 1, \eta + 1].$

The result of Proposition 7.4 can be understood as follows. By increasing β beyond $\eta + 1$, no additional collisions are prevented, but an increasing number of nodes is silenced. On the other hand, the nodes that become unblocked when decreasing β below $\eta - 1$, cause collisions when they activate. Note that for all values $\beta \in [\eta - 1, \eta + 1]$, we can rewrite (7.6) as

$$\theta(\beta,\eta,\sigma) = g(\beta) \cdot \frac{(\lambda_0(\beta))^{\beta-\eta-1}}{\beta+1},$$

with

$$g(\beta) = rac{\lambda_0(\beta) - 1}{\lambda_0(\beta) - rac{\beta}{\beta + 1}} \to 1, \quad \beta \to \infty.$$

We are now in the position to present our main result. While we already know that the optimal sensing range is contained in the interval $[\eta - 1, \eta + 1]$, the next result is more specific.

THEOREM 7.1. There exists a threshold interval $[\sigma_{\min}, \sigma_{\max}]$ such that

$$eta^* = \left\{ egin{array}{cc} \eta-1, & ext{if } \sigma \leq \sigma_{\min}, \ \eta+1, & ext{if } \sigma \geq \sigma_{\max}, \end{array}
ight.$$

and β^* increases from $\eta - 1$ to $\eta + 1$ when σ increases from σ_{\min} to σ_{\max} .

7.2 MAIN RESULTS

The proof of Theorem 7.1, see Appendix 7.A, follows from a detailed study of $\theta(\beta, \eta, \sigma)$ which involves implicit differentiation with respect to β (since $\lambda_0(\beta)$ is defined implicitly).

Theorem 7.1 can be interpreted as follows (see Figure 7.2). When σ is large, nodes activate very quickly after finishing their previous transmissions. When the system is in a maximal independent set, and if collisions are not ruled out, an activating node suffers a collision almost surely. This explains why for σ large, the optimal sensing range is $\beta = \eta + 1$, preventing collisions completely. On the other hand, when σ is small, collisions become rare, as few nodes are active simultaneously. In this case, the throughput is best served by increasing the spatial reuse, that is, decreasing the sensing range (up to $\eta - 1$). This explains the result of Theorem 7.1 for σ small.



Figure 7.2: The optimal sensing range β^* as a function of σ .

Note that Theorem 7.1 does not give the exact values of σ_{\min} and σ_{\max} . Instead, we give below an estimate of the location and width of the threshold interval.

THEOREM 7.2. Let $\kappa = \frac{\tau}{\eta+1}$ with $\tau = (\sqrt{5} - 1)/2$. (*i*) The threshold interval is bounded as

$$[\sigma_{\min}, \sigma_{\max}] \subseteq [\kappa(1+\kappa)^{\eta-1}, \kappa(1+\kappa)^{\eta+1}].$$

(ii) The width of the threshold interval is asymptotically given as

$$\sigma_{\max} - \sigma_{\min} \sim \frac{2e^{\tau}}{7+4\tau} \left(\frac{1}{\eta+1}\right)^2 \quad \text{as } \eta \to \infty.$$

Here we say that $f(\eta) \sim g(\eta)$ if $f(\eta)/g(\eta) \to 1$ as $\eta \to \infty$. From Theorem 7.2(ii) we see that the width of the threshold interval is $\mathcal{O}(\eta^{-2})$. Therefore, the interval width decreases rapidly as a function of η , and we can speak of an almost immediate transition from one regime ($\beta^* = \eta - 1$) to the other ($\beta^* = \eta + 1$). As a by-product of the proof of Theorem 7.2(ii) we obtain sharp approximations for σ_{\min} and σ_{\max} , see (7.31)-(7.32):

$$\hat{\sigma}_{\min} = \hat{\mu}_{-} (1 + \hat{\mu}_{-})^{\eta - 1}, \quad \hat{\sigma}_{\max} = \hat{\mu}_{+} (1 + \hat{\mu}_{+})^{\eta + 1},$$
(7.8)

with $\hat{\mu}_{\pm} = \tau / (\eta + \alpha_{\pm})$ and α_{\pm} given as α in (7.32) with $\gamma = \pm 1$.

OPTIMAL TRADEOFF BETWEEN EXPOSED AND HIDDEN NODES

7.2.1 Throughput limiting behavior

We now consider some limiting regimes for which we can make more explicit statements about the throughput. From Theorem 7.2 we can already see that the threshold interval moves in the direction of zero as η becomes large which implies that $\beta^* = \eta + 1$ for small values of σ . The next result shows that in the regime where η becomes large, the maximum throughput tends to zero.

PROPOSITION 7.5. Let $\sigma > 0$ be fixed. As $\eta \rightarrow \infty$,

$$\max_{\beta} \theta(\beta, \eta, \sigma) = \frac{1}{\eta + 2} \left(1 + \mathcal{O}\left(\frac{1}{\ln(\eta + 1)}\right) \right).$$

For $\beta \ge \eta + 1$ our model reduces to a model without collisions that was studied extensively in [6, 11, 22, 71, 106], as well as Chapters 4-6. In particular, one immediately obtains from (7.6) the following result:

COROLLARY 7.1. Let $\beta \ge \eta + 1$. Then

$$\theta(\beta,\eta,\sigma) = \frac{\lambda_0 - 1}{(\beta + 1)\lambda_0 - \beta}$$

This result was also derived in [6, 22, 71, 106]. From Proposition 7.7 and the proof of Proposition 7.5 it is seen that $\lambda_0 \to \infty$ as $\sigma \to \infty$ and β is fixed, and that $\beta(\lambda_0 - 1) \to \infty$ as $\beta \to \infty$ and σ is fixed. Thus the throughput is approximately $\frac{1}{\beta+1}$ when either σ or β is large. This can be understood as follows. For large σ , the high activity rate allows for configurations close to the maximum-size independent set: A configuration in which one out of every $\beta + 1$ nodes in active. For β large, when a node deactivates, a large number of neighboring nodes become eligible for activation. The time until the first such node activates goes to 0 when β increases.

COROLLARY 7.2. Let $\beta \leq \eta$. Then

$$\lim_{\sigma \to \infty} \theta_n(\beta, \eta, \sigma) = 0.$$
(7.9)

PROOF. From (7.17) it follows that

$$\lambda_0(\sigma) = \sigma^{\frac{1}{1+\beta}} + \mathcal{O}(1), \quad \sigma \to \infty.$$
(7.10)

Substituting (7.10) into (7.6), and using that $f(\beta) > 2\beta$ when $\beta \le \eta$, yields

$$\theta_n(\beta,\eta,\sigma) = \frac{\sigma(\sigma^{\frac{1}{1+\beta}} + \mathcal{O}(1))^{\beta - f(\beta)}}{(\beta + 1)(\sigma^{\frac{1}{1+\beta}} + \mathcal{O}(1)) - \beta} \to 0, \quad \sigma \to \infty,$$

which gives (7.9).

Figure 7.3 shows the throughput plotted against the activity rate σ for $\eta = 7$ and various values of β . When $\beta \leq \eta$, the throughput gradually drops to 0, whereas for $\beta \geq \eta + 1$, the throughput will eventually converge to the limit $1/(\beta + 1)$. This confirms Corollaries 7.1 and 7.2.

7.3 NORMALIZATION CONSTANT ROOTS



Figure 7.3: The throughput $\theta(\beta, \eta, \sigma)$ plotted against σ for $\eta = 7$ and various values of β .

7.3 Normalization constant roots

In this section we study the roots $\lambda_0, \ldots, \lambda_\beta$ of (7.3) in more detail. In particular, we derive exact infinite-series expressions for the roots that are used in this chapter both for numerical purposes (in Section 7.5) and to prove Corollary 7.2. These roots are essential in Section 7.4.1, where the finite and infinite networks are compared. Our main tool will be the Lagrange inversion theorem (see [16, p.22]), and depending on the value of σ , this gives two different infinite-series expressions. Let $(x)_n = \Gamma(x+n)/\Gamma(x)$ denote the Pochhammer symbol.

PROPOSITION 7.6. For small $\sigma > 0$,

$$\lambda_0(\sigma) = 1 + \sum_{l=1}^{\infty} \frac{(-1)^{l-1} (\beta l)_{l-1}}{l!} \sigma^l,$$
(7.11)

111

$$\lambda_j(\sigma) = \sum_{l=1}^{\infty} \frac{(l/\beta)_{l-1}}{l!} w_j^l, \quad j = 1, 2, \dots, \beta,$$
(7.12)

where $w_j = \sigma^{1/\beta} e^{2\pi i (j-1/2)/\beta}$ and $i = \sqrt{-1}$. The series expansions in (7.11) and (7.12) converge for

$$0 \le \sigma \le \frac{\beta^{\beta}}{(\beta+1)^{\beta+1}} =: \xi(\beta), \tag{7.13}$$

and diverge otherwise.

PROOF. We first consider the case j = 0. Set $\mu_0 = \lambda_0 - 1$, so μ_0 satisfies $\mu_0(1+\mu_0)^\beta = \sigma$. Hence for small values of $|\sigma|$ we have by Lagrange's inversion theorem

$$\mu_{0} = \sum_{l=1}^{\infty} \frac{1}{l!} \left(\frac{\mathrm{d}}{\mathrm{d}\mu} \right)^{l-1} \left[\left(\frac{\mu}{\mu(1+\mu)^{\beta}} \right)^{l} \right]_{\mu=0} \sigma^{l} = \sum_{l=1}^{\infty} \frac{(-1)^{l-1} (\beta l)_{l-1}}{l!} \sigma^{l}.$$
(7.14)

Next we consider the case that $j = 1, ..., \beta$. We now write (7.3) as

$$\lambda^{\beta}(1-\lambda) = -\sigma, \quad \lambda(1-\lambda)^{1/\beta} = w_j,$$

OPTIMAL TRADEOFF BETWEEN EXPOSED AND HIDDEN NODES

where

$$w_i = \sigma^{1/\beta} \mathrm{e}^{2\pi i (j-1/2)/\beta}$$

Then we get for $|w_j|$ sufficiently small

$$\lambda_j = \sum_{l=1}^{\infty} \frac{1}{l!} \left(\frac{\mathrm{d}}{\mathrm{d}\lambda}\right)^{l-1} \left[\left(\frac{\lambda}{\lambda(1-\lambda)^{1/\beta}}\right)^l \right]_{\lambda=0} w_j^l = \sum_{l=1}^{\infty} \frac{(l/\beta)_{l-1}}{l!} w_j^l.$$
(7.15)

The radii of convergence of the series in (7.14) and (7.15) are easily obtained from the asymptotics

$$\Gamma(x+1) = x^{x+1/2} e^{-x} \sqrt{2\pi} (1 + \mathcal{O}(x^{-1})), \quad x \to \infty,$$
(7.16)

of the Γ -function, used to examine the Pochhammer quantities $(x)_n = \Gamma(x + n)/\Gamma(x)$ and the factorials $l! = \Gamma(l + 1)$ that occur in both series. This yields the result that both series converge when $\sigma \leq \xi(\beta)$ and diverge for $\sigma > \xi(\beta)$. When $\sigma = \xi(\beta)$ the terms in either series are $\mathcal{O}(l^{-3/2})$.

PROPOSITION 7.7. For large $\sigma > 0$,

$$\lambda_{j}(\sigma) = \left(\sum_{l=1}^{\infty} \frac{\left(\frac{-l}{\beta+1}\right)_{l-1}}{l!} \nu_{j}^{-l}\right)^{-1}, \quad j = 0, 1, \dots, \beta,$$
(7.17)

where $v_i = \sigma^{1/(\beta+1)} e^{2\pi i j/(\beta+1)}$. The series expansion in (7.17) converges for

$$\sigma \geq \xi(\beta),$$

and diverges otherwise, where $\xi(\beta)$ is given in (7.13).

PROOF. We can treat the cases j = 0 and $j = 1, ..., \beta$ simultaneously now. We write (7.3) in the form

$$\frac{1}{\lambda} \left(1 - \frac{1}{\lambda}\right)^{\frac{-1}{\beta+1}} = \left(\frac{1}{\sigma}\right)^{\frac{1}{\beta+1}} = \nu^{-1},$$

where we let

$$\nu^{-1} = \nu_j^{-1} = \left(\frac{1}{\sigma}\right)^{\frac{1}{\beta+1}} e^{-2\pi \imath \frac{j}{\beta+1}}, \quad j = 0, 1, \dots, \beta$$
(7.18)

with $\sigma^{-\frac{1}{\beta+1}} > 0$ in (7.18). We get for sufficiently large σ from Lagrange's inversion theorem (with $u = 1/\lambda$) that

$$\frac{1}{\lambda_j} = \sum_{l=1}^{\infty} \frac{1}{l!} \left(\frac{\mathrm{d}}{\mathrm{d}u}\right)^{l-1} \left[\left(\frac{u}{u(1-u)^{-1/(\beta+1)}}\right)^l \right]_{u=0} v_j^{-l} = \sum_{l=1}^{\infty} \left(\frac{-l}{\beta+1}\right)_{l-1} \frac{v_j^{-l}}{l!}.$$
 (7.19)

The Pochhammer quantity $(\frac{-l}{\beta+1})_{l-1}$ vanishes if and only if l = 1, 2, ... is a multiple of $\beta + 1$. The radius of convergence of the series in (7.19) is again determined by the asymptotics of the Γ -function in (7.16). Here it must also be used that

$$\Gamma(-J) = \frac{-1}{\Gamma(J+1)} \frac{\pi}{\sin \pi J}, \quad J > 0$$

It follows that the series in (7.19) is convergent when $\sigma \ge \xi(\beta)$ and divergent when $\sigma < \xi(\beta)$. When $\sigma = \xi(\beta)$ the terms in the series are $\mathcal{O}(l^{-3/2})$.



Figure 7.4 shows the roots of (7.3) drawn in the complex λ -plane for $\beta = 4$. Each heavy solid line corresponds to a root as a function of σ , and the dots represent the threshold $|\sigma| = \xi(\beta)$. The light solid straight line and the dashed straight line illustrate the leading behavior of each root as $\sigma \downarrow 0$ or $\sigma \rightarrow \infty$ according to Propositions 7.6 and 7.7, respectively. The dashed curve encircling the origin 0 and the point 1 is the image of $\nu \in \mathbb{C}$ with $|\nu| = \sigma^{1/(\beta+1)}$, $\sigma = \xi(\beta)$, under the mapping given by the reciprocal of the right-hand side of (7.17) with ν_j replaced by ν .

113



Figure 7.4: The roots of $\lambda^{\beta+1} - \lambda^{\beta} = \sigma$ as functions of σ in (7.11), (7.12) and (7.17), for $\beta = 4$.

7.4 Optimal sensing range for general networks

We now discuss two remaining issues. In Section 7.4.1 we investigate to what extent the asymptotic results give accurate predictions for *finite* line networks. In Section 7.4.2 we investigate whether the notions of two regimes and a critical threshold carry over to more general topologies.

7.4.1 Finite versus infinite line networks

We now look at the approximation error $|\theta_n - \theta|$ and the resulting error in the optimal sensing range. To investigate the error we plot θ_n and θ in Figure 7.5, represented by the dashed line and the solid line, respectively. All results for θ_n were obtained by using (7.4) and (7.5) in combination with the infinite-series expressions for the roots in Section 7.3.



We take n = 100 (201 nodes), $\eta = 4$, and we let β increase from 1 to 100. In Figure 7.5(a) $\sigma = 0.25$, and in Figure 7.5(b) $\sigma = 5$. For β small the error $|\theta_n(\beta) - \theta(\beta)|$ is negligible, but the error increases as β increases. This can be explained by the observation that for larger β , the number of roots of (7.3) increases, as does the number of roots discarded by the approximation. This phenomenon becomes more pronounced for larger values of σ . The non-monotone behavior of θ_n is caused by the fact that for finite n, the system is directed to maximum-size independent sets of active nodes, in particular for σ large, and these sets change dramatically with β . The most important observation is that the error $|\theta_n - \theta|$ is small for those values of β that lead to a large throughput. Figure 7.6 is similar to Figure 7.5, but instead of



Figure 7.5: The throughput θ_n (dashed) and θ (solid) plotted against β (with n = 100).

fixing *n* and varying β , we set $\beta = 16$ and vary *n*. In Figure 7.6(a) we take $\sigma = 0.25$ and in Figure 7.6(b) we take $\sigma = 5$. The accuracy of the approximation increases with *n*.



Figure 7.6: The throughput θ_n (dashed) and θ (solid) plotted against *n* (with $\beta = 16$).

Figure 7.7 shows the optimal sensing range plotted against σ , for $\eta = 5$. Each of the Figures 7.7(a)-7.7(d) shows the optimal range $\beta_n^*(\sigma)$ for finite *n*. We take $\eta = 5$ for all figures, and let σ increase from 0.15 to 0.19. The vertical lines indicate the approximations of the threshold interval from (7.8), and we see that these are sharp. The optimal sensing range β^* for $n \to \infty$ behaves as predicted by Theorem 7.1,



jumping from $\eta - 1$ before the threshold interval, to $\eta + 1$ after this interval, and β_n^* shows a similar pattern. We conclude that $n \to \infty$ provides a good approximation for the behavior of finite-sized networks, already for small and moderate values of *n*.

115



Figure 7.7: The optimal sensing range β_n^* (dashed) and β^* (solid) plotted against σ around the threshold interval for various values of *n* and $\eta = 5$.

7.4.2 General topologies

To investigate more general topologies, we first need a more elaborate description of the model. In addition to nodes, we introduce directed links between nodes that represent the possibility of transmissions taking place between these nodes. For two nodes to be able to transmit data, we require them to be within (Euclidian) distance *m* of each other. We assume links are formed between all nodes within distance *m*. Each node has back-off rate σ , and the destination of a transmission is chosen uniformly from all links originating from the activating node. The sensing range β and interference range η are also defined using the Euclidian distance.

First we consider 16 nodes placed on a 4×4 grid at unit distance from each other. The grid is wrapped around (top and bottom nodes on any vertical line and left and right nodes on any horizontal line are connected) so that the network is fully symmetric and all nodes have the same environment (and the same throughput), eliminating boundary effects. We set m = 1 and construct links between neighboring nodes (see Figure 7.8(a)). We take $\eta = 1$ and $\beta = 0, 1, 1.5, 2$. We run a discrete-event simulation of the dynamics described above.



Figure 7.8(b) shows the average per-node throughput plotted against σ . For σ small we see that $\beta = 0$ (i.e. $\beta = \eta - m$) is throughput-optimal, and for σ large it turns out $\beta = 2$ ($\beta = \eta + m$) is optimal. Moreover, when β is such that collisions can occur ($\beta < 2$), we see that the throughput decreases when σ increases, while for $\beta = 2$ the throughput approaches a non-zero limiting value for large σ .



Figure 7.8: A grid network and the corresponding per-node throughput.

We next show in Figure 7.9 a randomly generated network with 16 nodes. The transmission ranges are indicated by the circles, and links are displayed as lines. We assume a transmission range of m = 1 and interference range $\eta = 1.6$. Links are formed between all nodes within distance m and when a node activates, it uniformly chooses a node within distance m as the receiver.



Figure 7.9: Random network with 16 nodes.

The simulation results for the network in 7.9 are shown in Figure 7.10. The average per-node throughput is plotted against σ for $\beta = 0.2, 0.3, 1, 1.3, 1.5$. Figure 7.10 shows resemblance with Figure 7.3 for the infinite line. For β small the throughput drops as σ increases, as a result of collisions. For large β collisions are precluded, and the

7.5 CONCLUDING REMARKS



Figure 7.10: The average per-node throughput plotted against σ .

average throughput stabilizes. Moreover, we see that the optimal sensing range β^* again depends on σ . For $\sigma < 0.1$ we have $\beta^* = 0.3$ (this is not visible in the picture), whereas for $\sigma > 0.1$ the optimal sensing range is $\beta^* = 1$.

The tradeoff for individual nodes in an irregular network is more complicated. Although we see a similar threshold interval ($\sigma_{\min}, \sigma_{\max}$) that separates two sensing regimes, the position of the threshold interval and the optimal sensing range may differ between nodes. This depends on the direct surroundings of the node, as well as on the entire network structure.

7.5 Concluding remarks

In this chapter we studied a linear CSMA network in the presence of collisions. We considered the problem of determining the carrier-sensing range that maximizes the throughput, which amounts to a tradeoff between hidden nodes and exposed nodes. In order to get a handle on the throughput function we studied the wireless network in the asymptotic regime of infinitely many nodes. This resulted in a tractable limiting expression for the throughput of node zero (and hence of any other node) that allowed us to prove the following two results:

(i) To optimize the throughput, one should always choose a sensing range β that is close to the interference range η , and in fact the optimal sensing range is contained in the interval $[\eta - 1, \eta + 1]$ (see Proposition 7.4).

(ii) The sensing range β^* that optimizes the throughput equals $\eta - 1$ for less aggressive nodes (small σ) and $\eta + 1$ for aggressive nodes (large σ). In fact, we were able to show the existence of a threshold interval for σ that separates these two regimes (Theorem 7.1). This result provides (partial) justification for the frequently made assumption that no collisions occur. Indeed, one key insight is that if σ is large enough, ruling out all collisions by setting $\beta = \eta + 1$ is optimal.

We have further shown that the threshold interval is in many cases small, which implies that one can speak of an almost immediate transition from one regime ($\beta^* = \eta - 1$) to the other ($\beta^* = \eta + 1$). We have argued that, when the aggressiveness of the

Optimal tradeoff between exposed and hidden nodes

nodes is large enough, the system no longer gains from the potential benefits of more flexibility (small β), and just settles for the situation with no collisions.

7.A REMAINING PROOFS

Appendix

7.A Remaining proofs

7.A.1 Proof of Proposition 7.1

We write the generating function from (7.2) as

$$Z(x,\sigma) = \frac{P(x)}{S(x)},$$

where

$$P(x) = 1 + \sigma \frac{x^{\beta+1} - x}{x - 1}, \quad S(x) = 1 - x - \sigma x^{\beta+1}.$$

It is shown in [71] that the equation S(x) = 0 has $\beta + 1$ roots x_j , $j = 0, 1, ..., \beta$, and exactly one of them, x_0 is real and positive, while $|x_j| > x_0$, $j = 1, ..., \beta$. To prove Proposition 7.1 we first need to establish that these roots are distinct.

PROPOSITION 7.8. The roots of S(x) = 0 are distinct.

PROOF. When S(x) = S'(x) = 0, we have

$$1 - x - \sigma x^{\beta + 1} = 0 = -1 - \sigma (\beta + 1) x^{\beta}.$$

This implies that $x = 1 + \frac{1}{\beta} > 1$ and so that $\sigma = \frac{1-x}{x^{\beta+1}} < 0$. However, σ is non-negative.

Now we proceed with the proof of Proposition 7.1. Let $\lambda_j = 1/x_j$ so that $\lambda = \lambda_j$ satisfies (7.3). Using that all zeros of *S* are distinct, we have for $Z(x, \sigma)$ the partial fraction expansion

$$Z(x,\sigma) = \sum_{j=0}^{\beta} \frac{P(x_j)}{S'(x_j)} \frac{1}{x - x_j}$$

Now

$$\frac{P(x_j)}{S'(x_j)} = \frac{1 + \sigma \frac{x_j^{\beta+1} - x_j}{x_j - 1}}{-1 - (\beta + 1)\sigma x_j^{\beta}} = \frac{-x_j^{-\beta}}{1 + (\beta + 1)\sigma x_j^{\beta}} = \frac{-x_j^{-\beta}}{1 + (\beta + 1)\frac{1 - x_j}{x_j}} = \frac{-\lambda_j^{\beta}}{(\beta + 1)\lambda_j - \beta}$$

Here it has been used that

$$\frac{1}{1-x_j} = \frac{-1}{\sigma x_j^{\beta+1}}, \quad \sigma x_j^{\beta} = \frac{1-x_j}{x_j}.$$

Then for $|x| < x_0$ we have

$$Z(x,\sigma) = \sum_{j=0}^{\beta} \frac{P(x_j)}{S'(x_j)} \sum_{i=0}^{\infty} \frac{-x^i}{x_j^{i+1}} = \sum_{i=0}^{\infty} x^i \left(\sum_{j=0}^{\beta} \frac{\lambda_j^{\beta+1}}{(\beta+1)\lambda_j - \beta} \lambda_j^i \right),$$

as required.

Optimal tradeoff between exposed and hidden nodes

7.A.2 Proof of Proposition 7.4

As introduced earlier,

$$\mu_0=\lambda_0-1.$$

Then μ_0 depends on β and σ , we have $\mu_0 > 0$, and

$$\mu_0 (1 + \mu_0)^\beta = \sigma. \tag{7.20}$$

a...

By implicit differentiation with respect to β , we get from (7.20) that

$$\frac{\partial\mu_0}{\partial\beta} = \frac{-\mu_0(1+\mu_0)\ln(1+\mu_0)}{1+\mu_0+\beta\mu_0}.$$
(7.21)

In particular, both μ_0 and λ_0 decrease as a function of $\beta > 0$. Consider the case that $0 \le \beta \le \eta - 1$. Using $\lambda_0^{\beta} = \frac{\sigma}{\lambda_0 - 1}$ we get

$$\theta(\beta,\eta,\sigma) = \sigma^2 \frac{\lambda_0^{-2\eta}}{(\lambda_0 - 1)((\beta + 1)\lambda_0 - \beta)} = \sigma^2 \frac{\lambda_0^{-2\eta}}{\mu_0(1 + \mu_0 + \beta\mu_0)}$$

Now $\lambda_0^{-2\eta}$ increases as a function of β , and we will show that $\mu_0(1+\mu_0+\beta\mu_0)$ decreases in $\beta > 0$. We have from (7.21) that

$$\begin{aligned} &\frac{\partial}{\partial\beta} \left[\mu_0 (1 + \mu_0 + \beta\mu_0) \right] = \frac{\partial}{\partial\beta} \left[\beta\mu_0^2 + \mu_0 + \mu_0^2 \right] \\ &= \mu_0^2 - \frac{1 + 2(1 + \beta)\mu_0}{1 + \mu_0 + \beta\mu_0} \mu_0 (1 + \mu_0) \ln(1 + \mu_0) \le \mu_0 (\mu_0 - (1 + \mu_0) \ln(1 + \mu_0)) < 0, \end{aligned}$$

where the last inequality follows from $x \ln x > x - 1, x > 1$. We conclude that θ increases as a function of $\beta \in (0, \eta - 1]$.

Next we consider the case that $\beta \ge \eta + 1$. From $\lambda_0^{\beta} = \frac{\sigma}{\lambda_0 - 1}$ we get

$$\theta(\beta,\eta,\sigma) = \sigma \frac{\lambda_0^{-\beta}}{(\beta+1)\lambda_0 - \beta} = \frac{\lambda_0 - 1}{(\beta+1)\lambda_0 - \beta} = \frac{\mu_0}{1 + \mu_0 + \beta\mu_0}.$$

Now

$$\frac{\partial}{\partial\beta}\left(\frac{\mu_0}{1+\mu_0+\beta\mu_0}\right)=\frac{\frac{\partial\mu_0}{\partial\beta}-\mu_0^2}{(1+\mu_0+\beta\mu_0)^2}<0,$$

see (7.21), and so θ decreases as a function of $\beta \ge \eta + 1$. Since θ depends continuously on $\beta > 0$, the result follows.

7.A REMAINING PROOFS

7.A.3 Proof of Theorem 7.1

The proof of the result as stated in Theorem 7.1 requires expanding several other results. We consider $\beta \in [\eta - 1, \eta + 1]$ so that

$$\theta(\beta,\eta,\sigma) = \sigma \frac{\lambda_0^{-\eta-1}}{(\beta+1)\lambda_0 - \beta} = \sigma \frac{(1+\mu_0)^{-\eta-1}}{1+\mu_0 + \beta\mu_0}$$

From (7.21) it follows from a straightforward but somewhat lengthy computation that

$$\frac{\partial}{\partial\beta} \left[\theta(\beta,\eta,\sigma)\right] = \frac{-\sigma\mu_0(1+\mu_0)^{-\eta-1}}{(1+\mu_0+\beta\mu_0)^2} \left(1-(\eta+2+\frac{\beta}{1+\mu_0+\beta\mu_0})\ln(1+\mu_0)\right).$$
(7.22)

Let

$$F(\beta,\sigma) = (\eta + 2 + \frac{\beta}{1 + \mu_0 + \beta\mu_0})\ln(1 + \mu_0).$$
(7.23)

Then we have for $\beta \in [\eta - 1, \eta + 1]$ that

$$F(\beta, \sigma) > 1 \Rightarrow \theta$$
 increases strictly in β , (7.24)

$$F(\beta, \sigma) < 1 \Rightarrow \theta$$
 decreases strictly in β . (7.25)

We analyze $F(\beta, \sigma)$ in some detail, especially for values of β, σ such that $F(\beta, \sigma) = 1$. We recall here that $\mu_0 = \mu_0(\beta, \sigma)$ is a function of β and σ as well.

We fix $\beta > 0$, and we compute

$$\frac{\partial}{\partial\beta}F(\beta,\sigma) = \left[\frac{\eta+1}{\mu_0+1} + \frac{1+\beta}{1+\mu_0+\beta\mu_0} - \frac{\beta(1+\beta)\ln(1+\mu_0)}{(1+\mu_0+\beta\mu_0)^2}\right]\frac{\partial\mu_0}{\partial\sigma}.$$

We get from (7.20) by implicit differentiation that

$$\frac{\partial \mu_0}{\partial \sigma} = \frac{\mu_0 (1 + \mu_0)}{\sigma (1 + \mu_0 + \beta \mu_0)} > 0.$$
(7.26)

Furthermore, it is seen from (7.20) that $\mu_0(\beta, \sigma) \to 0$ as $\sigma \downarrow 0$ and that $\mu_0(\beta, \sigma) \to \infty$ as $\sigma \to \infty$. Hence, $\mu_0(\beta, \sigma)$ increases from 0 to ∞ as σ increases from 0 to ∞ . Moreover,

$$\frac{\eta+1}{\mu_0+1} > 0, \quad 1 > \frac{\beta \ln(1+\mu_0)}{1+\mu_0+\beta\mu_0}.$$
(7.27)

It follows from (7.26) and (7.27) that $\frac{\partial}{\partial \sigma} F(\beta, \sigma) > 0$. Then, from (7.23) and from the fact that μ_0 increases from 0 to ∞ as σ increases from 0 to ∞ , we have that $F(\beta, \sigma)$ increases from 0 to ∞ as σ increases from 0 to ∞ . Therefore, for any $\beta > 0$, there is a unique $\sigma = \sigma(\beta)$ such that

$$F(\beta, \sigma) = F(\beta, \sigma(\beta)) = 1. \tag{7.28}$$

We will next show that $\sigma(\beta)$ increases in $\beta \in [\eta - 1, \eta + 1]$. By implicit differentiation in (7.28), we have for $\beta \in [\eta - 1, \eta + 1]$

$$0 = \frac{\mathrm{d}}{\mathrm{d}\beta} [F(\beta, \sigma(\beta))] = F_{\beta}(\beta, \sigma(\beta)) + \sigma'(\beta)F_{\sigma}(\beta, \sigma(\beta)), \qquad (7.29)$$

OPTIMAL TRADEOFF BETWEEN EXPOSED AND HIDDEN NODES

where F_{β} and F_{σ} denote the respective partial derivatives (and $\sigma'(\eta \pm 1)$ is the left and right derivative for + and –, respectively). We already know that $F_{\sigma} > 0$, and we will show now that $F_{\beta}(\beta, \sigma(\beta)) < 0$. To that end, we compute, using definition (7.23) of *F* and (7.21) that

$$\frac{\partial}{\partial \beta} [F(\beta,\sigma)] = -\ln(1+\mu_0) \Big[(\eta+2+\frac{\beta}{1+\mu_0+\beta\mu_0}) \frac{\mu_0}{1+\mu_0+\beta\mu_0} - \frac{1+\mu_0-\beta(1+\beta)\frac{\partial\mu_0}{\partial\beta}}{(1+\mu_0+\beta\mu_0)^2} \Big].$$

Next, from (7.23) and (7.28) we have that

$$\mu_0 \ge \ln(1 + \mu_0) = \frac{1}{\eta + 2 + \frac{\beta}{1 + \mu_0 + \beta \mu_0}}$$

and so

$$\begin{split} &\frac{\partial F}{\partial \beta}(\beta,\sigma(\beta)) \leq -\ln(1+\mu_0) \left[\frac{1}{1+\mu_0+\beta\mu_0} - \frac{1+\mu_0-\beta(1+\beta)\frac{\partial\mu_0}{\partial\beta}}{(1+\mu_0+\beta\mu_0)^2} \right]_{\sigma=\sigma(\beta)} \\ &= \frac{-\beta\ln(1+\mu_0)}{(1+\mu_0+\beta\mu_0)^2} \left[\mu_0 + (1+\beta)\frac{\partial\mu_0}{\partial\beta} \right]_{\sigma=\sigma(\beta)} \\ &= \frac{-\mu_0\beta\ln(1+\mu_0)}{(1+\mu_0+\beta\mu_0)^2} \left[1-(1+\beta)\frac{(1+\mu_0)\ln(1+\mu_0)}{1+\mu_0+\beta\mu_0} \right]_{\sigma=\sigma(\beta)}, \end{split}$$

where (7.21) has been used once more. Finally, from (7.23) and (7.28),

$$(1+\beta)\frac{(1+\mu_0)\ln(1+\mu_0)}{1+\mu_0+\beta\mu_0}\Big|_{\sigma=\sigma(\beta)} = \frac{(1+\beta)(1+\mu_0)}{(\eta+2)(1+\mu_0+\beta\mu_0)+\beta}\Big|_{\sigma=\sigma(\beta)} < 1,$$

since $0 < \beta \le \eta + 1$ and $\mu_0 > 0$. Hence, $F_{\beta}(\beta, \sigma(\beta)) < 0$ as required. It now follows from (7.29) and from $F_{\sigma}(\beta, \sigma(\beta)) > 0$ that $\sigma'(\beta) > 0$ when $\beta \in [\eta - 1, \eta + 1]$.

We have now shown that $\sigma(\beta)$ increases in $\beta \in [\eta - 1, \eta + 1]$. Next we let

 $\sigma_{\min} = \sigma(\eta - 1) < \sigma(\eta + 1) =: \sigma_{\max}.$

For $\sigma \in [\sigma_{\min}, \sigma_{\max}]$ the inverse function $\beta(\sigma) \in [\eta - 1, \eta + 1]$ increases in σ . It follows then from

$$F(\beta(\sigma), \sigma) = 1, \quad F_{\beta}(\beta(\sigma), \sigma) < 0$$

and (7.22)-(7.25) that $\theta(\beta, \eta, \sigma)$ is maximal at $\beta = \beta(\sigma)$ when $\sigma \in [\sigma_{\min}, \sigma_{\max}]$.

We will now complete the proof of Theorem 7.1. Let $\beta \in [\sigma_{\min}, \sigma_{\max}]$, and assume that $\sigma \leq \sigma_{\min}$. Then $\sigma < \sigma(\beta)$ and so $F(\beta, \sigma) < F(\beta, \sigma(\beta)) = 1$ since *F* increases in σ . Hence, θ strictly decreases at β . Similarly, θ strictly increases at $\beta \in (\eta - 1, \eta + 1)$ when $\sigma \geq \sigma_{\max}$. It follows that θ strictly decreases in $\beta \in [\eta - 1, \eta + 1]$ when $\sigma \leq \sigma_{\min}$ and that θ strictly increases in $\beta \in [\eta - 1, \eta + 1]$ when $\sigma \geq \sigma_{\max}$. Finally, when $\sigma \in (\sigma_{\min}, \sigma_{\max})$, we have that

$$F(\eta - 1, \sigma) > F(\eta - 1, \sigma_{\min}) = 1 = F(\eta + 1, \sigma_{\max}) > F(\eta + 1, \sigma)$$

showing that θ strictly increases at $\beta = \eta - 1$ and strictly decreases at $\beta = \eta + 1$, and assumes its maximum at $\beta = \beta(\sigma)$.

7.A REMAINING PROOFS

7.A.4 Proof of Theorem 7.2

We shall show below that

$$(\eta + 2 + \frac{\eta - 1}{1 + \eta\kappa})\ln(1 + \kappa) < 1 < (\eta + 2 + \frac{\eta + 1}{1 + (\eta + 2)\kappa})\ln(1 + \kappa)$$
(7.30)

where $\kappa = \tau/(\eta + 1)$. Assuming this, we recall that (for fixed $\beta > 0$) μ_0 strictly increases in σ and vice versa. Set

$$\sigma_{-} = \kappa (1+\kappa)^{\eta-1},$$

then $\kappa = \mu_0(\beta = \eta - 1, \sigma_-)$ and we have that $F(\eta - 1, \sigma_-) < 1$. So $\sigma_- < \sigma_{\min}$ since *F* is increasing in σ . Similarly, when

$$\sigma_+ = \kappa (1+\kappa)^{\eta+1},$$

we have that $\kappa = \mu_0(\beta = \eta + 1, \sigma_+)$ and then from (7.30) that $F(\eta + 1, \sigma_+) > 1$ and so $\sigma_+ > \sigma_{max}$. Therefore,

$$\sigma_{\max} - \sigma_{\min} < \sigma_{+} - \sigma_{-} = \kappa (1+\kappa)^{\eta-1} ((1+\kappa)^{2} - 1)$$

= $2\left(1 + \frac{\tau}{\eta+1}\right)^{\eta-1} \left(\frac{\tau}{\eta+1}\right) \left(1 + \frac{\tau}{\eta+1}\right) \le 2e^{\tau} \left(\frac{\tau}{\eta+1}\right)^{2} (1 + \frac{\tau}{\eta+1}).$

This proves Theorem 7.2(i). It remains to show (7.30). As to the first inequality in (7.30) we have

$$1 - (\eta + 2 + \frac{\eta - 1}{1 + \eta \kappa}) \ln(1 + \kappa) > 1 - (\eta + 2 + \frac{\eta - 1}{1 + \eta \kappa})\kappa$$
$$> \frac{1}{1 + \eta \kappa} (1 - (\eta + 1)\kappa - ((\eta + 1)\kappa)^2) = 0$$

since $1 - \tau - \tau^2 = 0$ and $(\eta + 1)\kappa = \tau$. As to the second inequality of (7.30) we have

$$\begin{split} &1 - (\eta + 2 + \frac{\eta + 1}{1 + (\eta + 2)\kappa})\ln(1 + \kappa) < 1 - (\eta + 2 + \frac{\eta + 1}{1 + (\eta + 2)\kappa})(\kappa - \frac{1}{2}\kappa^2) \\ &= \frac{1}{1 + (\eta + 2)\kappa} \Big(1 - (\eta + 1)\kappa - ((\eta + 1)\kappa)^2 - \kappa^2(\eta + 3/2 - \frac{1}{2}(\eta + 2)^2\kappa)\Big). \end{split}$$

As before

$$1 - (\eta + 1)\kappa - ((\eta + 1)\kappa)^2 = 0$$

and

$$\eta + \frac{3}{2} - \frac{1}{2}(\eta + 2)^2 \kappa = \eta + \frac{3}{2} - \frac{(\eta + 2)^2}{2(\eta + 1)}\tau > 0, \quad \eta \ge 0$$

since $\tau = \frac{1}{2}(\sqrt{5}-1) < \frac{3}{4}$ (which is the minimum value of $2(\eta + 3/2)(\eta + 1)(\eta + 2)^{-2}$ for $\eta \ge 0$). This shows the second inequality in (7.30).

We next prove Theorem 7.2(ii), and for this we need the following result:

Optimal tradeoff between exposed and hidden nodes

PROPOSITION 7.9. With $\beta = \eta + \gamma$ where $-1 \le \gamma \le 1$,

$$\sigma(\beta) = \mu(1+\mu)^{\eta+\gamma},\tag{7.31}$$

where

$$\mu = \frac{\tau}{\eta + \alpha + \mathcal{O}(\eta^{-1})}, \quad \alpha = \frac{(5 + 2\gamma)\tau + 1}{2(2\tau + 1)}, \quad (7.32)$$

and the \mathcal{O} holds uniformly in $\gamma \in [-1, 1]$.

PROOF. We have $\sigma(\beta) = \mu (1 + \mu)^{\beta}$ where μ is the unique solution of the equation

$$(\eta + 2 + \frac{\beta}{1 + (1 + \beta)\mu})\ln(1 + \mu) = 1.$$
(7.33)

We know from the proof of Theorem 7.2(i) that $\mu = O(\eta^{-1})$. Multiplying (7.33) by $1 + (1 + \beta)\mu$ and expanding

$$\ln(1+\mu) = \mu - \frac{1}{2}\mu^2 + \mathcal{O}(\mu^3),$$

we get

$$(\eta\beta + \frac{1}{2}\eta + \frac{3}{2}\beta + 1)\mu^{2} + (\eta + 1)\mu - 1 = \frac{1}{2}(\eta + 2)(\beta + 1)\mu^{3} + \mathcal{O}(\eta^{-2}).$$

Next let $\alpha \in \mathbb{R}$ be independent of η and use $\beta = \eta + \gamma$ to write

$$\eta\beta + \frac{1}{2}\eta + \frac{3}{2}\beta + 1 = (\eta + \alpha)^2 + (2 + \gamma - 2\alpha)\eta + \frac{3}{2}\gamma + 1 - \alpha^2.$$

Together with $\eta + 1 = \eta + \alpha + 1 - \alpha$, we obtain

$$\begin{aligned} &(\eta + \alpha)^2 \mu^2 + (\eta + \alpha)\mu - 1 \\ &= \frac{1}{2}(\eta + 2)(\eta + \gamma + 1)\mu^3 - ((2 + \gamma - 2\alpha)\eta + \frac{3}{2}\gamma + 1 - \alpha^2)\mu^2 - (1 - \alpha)\mu + \mathcal{O}(\eta^{-2}). \end{aligned} \tag{7.34}$$

We now take α such that the whole second term in of (7.34) is $\mathcal{O}(\eta^{-2})$. Using that $\mu = \frac{\tau}{n} + \mathcal{O}(\eta^{-2})$, this leads to

$$\frac{1}{2}\tau^{3} - (2 + \gamma - 2\alpha)\tau^{2} - (1 - \alpha)\tau = 0,$$

and this yields the α in (7.32). The polynomial $x^2 + x - 1 = 0$ has a zero of first order at $x = \tau$. Hence with α as in (7.32) we see from $(\eta + \alpha)^2 \mu^2 + (\eta + \alpha)\mu - 1 = \mathcal{O}(\eta^{-2})$ that $(\eta + \alpha)\mu = \tau + \mathcal{O}(\eta^{-2})$. This gives the result.

Now we proceed to prove Theorem 7.2(ii). We use the result of Proposition 7.9. Thus

$$\begin{split} \sigma(\eta+\gamma) &= \mu(1+\mu)^{\eta+\gamma}, \\ \mu &= \frac{\tau}{\eta+\alpha+\mathcal{O}(\eta^{-1})} = \frac{\tau}{\eta+\alpha}(1+\mathcal{O}(\eta^{-2})). \end{split}$$

7.A REMAINING PROOFS

By elementary considerations

$$\begin{aligned} \sigma(\eta + \gamma) &= \frac{\tau}{\eta + \alpha} (1 + \frac{\tau}{\eta + \alpha})^{\eta + \gamma} (1 + \mathcal{O}(\eta^{-2})) \\ &= \frac{\tau}{\eta + \alpha} \exp[(\eta + \gamma)(\frac{\tau}{\eta + \alpha} - \frac{\tau^2}{2(\eta + \alpha)})](1 + \mathcal{O}(\eta^{-2})) \\ &= \frac{\tau e^{\tau}}{\eta + \alpha} (1 + \frac{(\gamma - \alpha)\tau - \frac{1}{2}\tau^2}{\eta})(1 + \mathcal{O}(\eta^{-2})). \end{aligned}$$

Then letting $\gamma = \pm 1$ and

$$\alpha(1) = \frac{7\tau + 1}{2(2\tau + 1)}, \quad \alpha(-1) = \frac{3\tau + 1}{2(2\tau + 1)}$$

in accordance with Proposition 7.9, it follows that

$$\begin{aligned} \sigma(\eta+1) - \sigma(\eta-1) &= \frac{\tau e^{\tau}}{\eta^2} \Big(\alpha(-1) - \alpha(1) + (1 - \alpha(1))\tau + (1 + \alpha(-1))\tau \Big) + \mathcal{O}(\eta^{-3}) \\ &= \frac{\tau e^{\tau}}{\eta^2} \frac{2\tau^2}{2\tau + 1} + \mathcal{O}(\eta^{-3}). \end{aligned}$$

Finally, it follows easily from $\tau^2 + \tau = 1$ that $\tau^3(7 + 4\tau) = 2\tau + 1$.

7.A.5 Proof of Proposition 7.5

Since $\sigma > 0$ is fixed, it follows from (see the proof of Theorem 7.2)

$$\sigma_{\max} < \sigma_{+} = \frac{\tau}{\eta + 1} \left(1 + \frac{\tau}{\eta + 1} \right)^{\eta + 1} < \frac{\tau e^{\tau}}{\eta + 1}$$

that $\sigma_{\max} < \sigma$ when η is large enough. Then by Theorem 7.1

$$\max \theta = \theta(\eta + 1) = \frac{\lambda_0 - 1}{(\eta + 2)\lambda_0 - \eta - 1} = \frac{\mu_0}{(\eta + 2)\mu_0 + 1} = \frac{1}{\eta + 2} \frac{1}{1 + \frac{1}{(\eta + 2)\mu_0}}$$

where μ_0 is the unique positive real μ root of $\mu(1 + \mu)^{\eta+1} = \sigma$. We shall show that

$$(\eta+2)\mu_0 \ge \ln\sigma,\tag{7.35}$$

$$(\eta + 2)\mu_0 = \ln(\eta + 1) + \mathcal{O}(\ln\ln(\eta + 1)), \quad \eta \to \infty, \tag{7.36}$$

uniformly in $\sigma \in [\epsilon, M]$, where $\epsilon > 0$ and $M > \epsilon$ are fixed. To show (7.35), we note from $\mu_0(1 + \mu_0)^{\eta+1} = \sigma$ that

$$(\eta + 1)\mu_0 \ge (\eta + 1)\ln(1 + \mu_0) = \ln \sigma - \ln \mu_0. \tag{7.37}$$

Next $\sigma = \mu_0 (1 + \mu_0)^{\eta+1} \ge \mu_0^{\eta+2}$, and so $\ln \mu_0 \le \frac{1}{\eta+2} \ln \sigma$. Therefore

$$(\eta + 1)\mu_0 \ge \ln \sigma - \frac{1}{\eta + 2}\ln \sigma = \frac{\eta + 1}{\eta + 2}\ln \sigma,$$

and (7.35) follows. As to (7.36), we first observe from (7.21) that μ_0 decreases in η when $\sigma > 0$ is fixed. Hence $L = \lim_{\eta \to \infty} \mu_0$ exists, and it follows from $\mu_0 (1+\mu_0)^{\eta+1} = \sigma$

Optimal tradeoff between exposed and hidden nodes

that L = 0. Thus, μ_0 decreases to 0 as $\eta \to \infty$. Then, from (7.37) we get that $(\eta + 1)\mu_0$ increases to ∞ as $\eta \to \infty$. All this holds uniformly in $\sigma \in [\epsilon, M]$: Since μ_0 increases in σ , the right-hand side of (7.37) is bounded below by $\ln \epsilon - \ln \mu_0(\sigma = M)$. Now take $\eta_0 > 0$ such that $(\eta + 1)\mu_0 \ge \sigma$ when $\eta \ge \eta_0$ and $\epsilon \le \sigma \le M$. Then from $\mu_0(1 + \mu_0)^{\eta + 1} = \sigma$ we have

$$(\eta + 1)\ln(1 + \mu_0) = \ln \sigma - \ln \mu_0 \le \ln(\eta + 1)\mu_0 - \ln \mu_0 \le \ln(\eta + 1)$$

when $\eta \ge \eta_0$ and $\epsilon \le \sigma \le M$. Hence, when $\eta \ge \eta_0$,

$$\mu_0 \le \exp\left[\frac{\ln(\eta+1)}{\eta+1}\right] - 1 = \frac{\ln(\eta+1)}{\eta+1} + \mathcal{O}\left(\left(\frac{\ln(\eta+1)}{\eta+1}\right)^2\right),\tag{7.38}$$

where the \mathcal{O} holds uniformly in $\sigma \in [\epsilon, M]$. Then, by (7.37),

$$(\eta + 1)\mu_{0}$$

$$\geq \ln \sigma - \ln \left(\exp \left[\frac{\ln(\eta + 1)}{\eta + 1} \right] - 1 \right) = \ln \sigma - \ln(\frac{\ln(\eta + 1)}{\eta + 1} \left(1 + \mathcal{O}\left(\frac{\ln(\eta + 1)}{\eta + 1} \right) \right)$$

$$= \ln(\eta + 1) - \ln \ln(\eta + 1) + \ln \sigma + \mathcal{O}\left(\frac{\ln(\eta + 1)}{\eta + 1} \right), \quad (7.39)$$

with \mathcal{O} holding uniformly in $\sigma \in [\epsilon, M]$ and $\eta \ge \eta_0$. From (7.38) and (7.39) we get (7.35) uniformly in $\sigma \in [\epsilon, M]$.

8 Time-slotted CSMA

In this chapter we study the performance of a time-slotted CSMA algorithm, where nodes are completely synchronized and transmissions last one time slot. The performance measures of interest are the same as for the continuous-time CSMA model: throughput, fairness and stability. We first look at the throughput under saturation assumptions, and compute the network-aggregate throughput as well as the per-node throughputs. The latter can be used to study fairness, similar to Chapter 5 in the case of continuous-time CSMA. We then relax the saturation assumption and consider a multi-hop network, in which packets are forwarded through the network. We study the stability of each node, and derive bounds on the end-to-end throughput.

The present model is different from the continuous-time CSMA model introduced in Section 1.3.2 and studied in Chapters 4-7, as these chapters assume that nodes operate asynchronously and in continuous time. As in Chapters 5 and 7, we consider a linear network. The multi-hop network discussed in Sections 8.3 and 8.4 is similar to the network presented in Section 5.5.

This chapter is structured as follows. In Section 8.1 we describe slotted CSMA and introduce the model of interest. Section 8.2 is devoted to throughput and fairness in the saturated case, while the unsaturated model is introduced and analyzed in Section 8.3. In Section 8.4 we compare the performance of slotted CSMA and continuous-time CSMA, and Section 8.5 offers some concluding remarks.

TIME-SLOTTED CSMA

8.1 Model description

We consider a linear network of *n* nodes which can be either active or inactive, depending on whether they are transmitting or not. Nodes within distance β are prevented from simultaneous activity. Similar to the CSMA model discussed in Chapters 4-7, the state of the network can be written as

$$\boldsymbol{\omega} = (\omega_1, \omega_2, \dots, \omega_n) \in \{0, 1\}^n,$$

where $\omega_i = 1$ when node *i* is active. The set of feasible states $\Omega \subseteq \{0, 1\}^n$ is the same as for the linear networks in Chapters 5 and 7, i.e., $\omega \in \{0, 1\}^n$ is feasible if and only if $\omega_i \omega_j = 0$ for all *i*, *j* such that $1 \le |i - j| \le \beta$. We assume that each node is saturated, i.e., it always has packets available for transmission. This assumption is relaxed in Section 8.3.

Time is slotted, and at the beginning of each time slot a feasible subset of nodes is activated for the duration of that slot. We denote the schedule of slot t by $\mathbf{X}(t) = (X_1(t), X_2(t), \ldots, X_n(t)) \in \Omega$, with $X_i(t) = 1$ if node i is active in slot t and $X_i(t) = 0$ otherwise. The states $\mathbf{X}(t)$ are i.i.d. across time, and are generated as follows. At the beginning of each time slot a random permutation $\mathbf{A}(t) = (A_1(t), A_2(t), \ldots, A_n(t))$ is chosen uniformly from the set of all n! permutations. Here $A_i(t)$ denotes the index of node i in the permutation $\mathbf{A}(t)$. Nodes then activate according to

$$X_i(t) = \begin{cases} 1, & \text{if } X_j(t) = 0 \ \forall j : A_j(t) < A_i(t) \text{ and } |i-j| \le \beta, \\ 0, & \text{otherwise,} \end{cases}$$
(8.1)

starting from the node with the lowest index.

Thus nodes activate in the order prescribed by the permutation **A**, but only if no other nodes within distance β are already active. This procedure yields for each time slot a feasible state $\omega \in \Omega$. A closed-form expression for the distribution of the activity process $\mathbb{P}(\mathbf{X}(t) = \omega)$ remains elusive, in contrast to the CSMA model (see (1.8)). However, as we will see in this chapter, we do not require such distribution in order to study the throughput.

The above procedure can be implemented in a distributed fashion by synchronizing all nodes and partitioning each time slot in a contention period and a data period. At the beginning of a contention period all nodes draw a uniformly distributed backoff time between 0 and the length of the contention period. A node activates when its back-off timer runs out, but only if no nodes within distance β are already active. Nodes then transmit for the entire duration of the data period. The duration of the contention period has to be sufficiently large to allow the carrier-sensing mechanism to function correctly. However, it can always be assumed to be much smaller than the length of the data period by scaling up the transmission durations. In the remainder of this chapter we assume the length of the contention period to be zero, and we arrive at the algorithm in (8.1).

Synchronization has long been part of the IEEE 802.11 protocol in the case of small networks where all nodes can communicate directly with each other. The recent 802.11s (mesh) amendment also provides synchronization for large networks.

8.2 The saturated regime

8.2 The saturated regime

Recall that the throughput of a node is defined as the rate at which successful transmissions are completed. For slotted CSMA, this is equivalent to the fraction of slots a node is active. Let us denote by $T_i(n)$ the throughput of node *i* in an *n*-node network, and by $E_n = \sum_{i=1}^n T_i(n)$ the network-aggregate throughput.

Throughout this section we will restrict ourselves to the case $\beta = 1$. The following proposition presents the aggregate throughput in this case.

PROPOSITION 8.1. The aggregate throughput in a network of n nodes is given by

$$E_n = \sum_{k=1}^n (-1)^{k+1} \frac{2^{k-1}}{k!} (n-k+1).$$

PROOF. Conditioning on the position of the first node to activate, we may write

$$E_{n} = 1 + \frac{2}{n}E_{n-2} + \sum_{k=2}^{n-1} \frac{1}{n}(E_{k-2} + E_{n-k-1}) = 1 + \sum_{k=1}^{n} \frac{1}{n}(E_{k-2} + E_{n-k-1})$$
$$= 1 + 2\sum_{k=1}^{n} \frac{1}{n}E_{k-2} = 1 + \frac{2}{n}\sum_{k=1}^{n} E_{k-2},$$
(8.2)

with the convention that $E_0 = E_{-1} = 0$. The generating function

$$\phi(\rho) = \sum_{n=1}^{\infty} E_n \rho^n, \qquad (8.3)$$

is well defined for any $0 \le \rho < 1$, since $0 \le E_n \le n$. In order to determine $\phi(\rho)$, we compute

$$\phi'(\rho) = \sum_{n=1}^{\infty} nE_n \rho^{n-1} = \sum_{n=1}^{\infty} n\left(1 + \frac{2}{n}\sum_{k=1}^n E_{k-2}\right)\rho^{n-1}$$
$$= \frac{1}{(1-\rho)^2} + \frac{2}{1-\rho}\sum_{k=1}^{\infty} E_{k-2}\rho^{k-1} = \frac{1}{(1-\rho)^2} + \frac{2\rho}{1-\rho}\phi(\rho).$$
(8.4)

The system in (8.4) is a standard first-order differential equation, with initial condition $\phi(0) = 0$, so that

$$\phi(\rho) = \frac{1 - e^{-2\rho}}{2(1 - \rho)^2}.$$

Now use

$$1 - e^{-2\rho} = \sum_{n=1}^{\infty} (-1)^{n+1} \frac{2^n}{n!} \rho^n \text{ and } (1 - \rho)^{-2} = \sum_{n=1}^{\infty} n\rho^{n-1} = \sum_{n=0}^{\infty} (n+1)\rho^n$$

to conclude that

$$\phi(\rho) = \frac{1}{2} \sum_{n=1}^{\infty} \left(\sum_{k=1}^{n} (-1)^{k+1} \frac{2^k}{k!} (n-k+1) \right) \rho^n,$$

completing the proof.

TIME-SLOTTED CSMA

We now turn to the individual throughputs in order to gain a more detailed understanding of the network. The following theorem gives a closed-form expression for the per-node throughputs $T_i(n)$, i = 1, 2, ..., n.

THEOREM 8.1. For $n \ge 1$ and $1 \le i \le n$,

$$T_{i}(n) = \begin{cases} \sum_{k=0}^{n-i} d_{i,i+2k}, & \text{if } (n-i) \text{ is even,} \\ \sum_{k=0}^{n-i-1} d_{i,i+2k+1}, & \text{if } (n-i) \text{ is odd,} \end{cases}$$
(8.5)

where

$$d_{i,n} = -a_i \frac{(-1)^{n-i}}{(n-i)!} + (-1)^i b_{i,n},$$
(8.6)

and

$$a_{i} = \sum_{k=0}^{i-1} \frac{(-1)^{k}}{k!}, \quad b_{i,n} = \sum_{k=0}^{i-1} \frac{(-1)^{n-k}}{k!(n-k)!}.$$
(8.7)

PROOF. Conditioning on the first node to activate yields the following recursive equation:

$$T_{i}(n) = \frac{1}{n} + \frac{1}{n} \sum_{j=1}^{i-2} T_{i-j-1}(n-j-1) + \frac{1}{n} \sum_{j=i}^{n-2} T_{i}(j).$$
(8.8)

With $\psi_i(\rho) = \sum_{n=i}^{\infty} T_i(n)\rho^n$, summing (8.8) over *n* gives the differential equation

$$\psi_i'(\rho) = \sum_{j=1}^{i-2} \rho^j \psi_{i-j-1}(\rho) + \frac{\rho^{i-1}}{1-\rho} + \frac{\rho}{1-\rho} \psi_i(\rho), \tag{8.9}$$

with initial condition $\psi_i(0) = 0$.

We shall show below that

$$\psi_i(\rho) = \frac{1}{1-\rho^2} \Big(\rho^i + (-1)^{i+1} - e^{-\rho} \rho^i a_i + (-1)^i e^{-\rho} \sum_{k=0}^{i-1} \frac{\rho^k}{k!} \Big), \tag{8.10}$$

which leaves (8.5) to prove. To this end we shall find the Taylor expansion for (8.10) with respect to the powers of ρ . Let us start with the last term inside the brackets in (8.10):

$$e^{-\rho} \sum_{k=0}^{i-1} \frac{\rho^k}{k!} = \sum_{m=0}^{\infty} \frac{(-1)^m \rho^m}{m!} \cdot \sum_{k=0}^{i-1} \frac{\rho^k}{k!} = \sum_{m=0}^{\infty} c_m \rho^m$$

with

$$c_m = \begin{cases} 1, & \text{if } m = 0, \\ \sum_{k=0}^{n} \frac{(-1)^{m-k}}{k!(m-k)!} = 0, & \text{if } 0 < m \le i-1, \\ \sum_{k=0}^{i-1} \frac{(-1)^{m-k}}{k!(m-k)!} = b_{i,m}, & \text{if } m \ge i, \end{cases}$$

8.2 THE SATURATED REGIME

and $b_{i,m}$ as in (8.7). Substituting this into (8.10) and using the Taylor expansion for the exponential function

$$e^{-\rho s} = \sum_{m=0}^{\infty} (-1)^m \frac{\rho^m}{m!} s^m$$
(8.11)

131

yields

$$\begin{split} \psi_i(\rho) &= \frac{1}{1-\rho^2} \Big(\rho^i + a_i \sum_{m=i}^{\infty} \frac{(-1)^{m-i}}{(m-i)!} \rho^m + (-1)^i \sum_{m=i}^{\infty} b_{i,m} \rho^m \Big) \\ &= \frac{1}{1-\rho^2} \Big(\rho^i + \sum_{m=i}^{\infty} d_{i,m} \rho^m \Big), \end{split}$$

with $d_{i,m}$ defined in (8.6). The $T_i(n)$ then readily follow from $\psi_i(\rho)$.

This proves Theorem 8.1. It remains to be shown that (8.10) holds. Introducing

$$\nu(\rho,s) = \sum_{i=1}^{\infty} \psi_i(\rho) s^i,$$

and using (8.9) gives

$$\frac{\partial v(\rho, s)}{\partial \rho} = \sum_{i=1}^{\infty} \psi_i'(\rho) s^i = \sum_{i=1}^{\infty} \sum_{j=1}^{i-2} \rho^j \psi_{i-j-1}(\rho) s^i + \sum_{i=1}^{\infty} \frac{\rho^{i-1} s^i}{1-\rho} + \sum_{i=1}^{\infty} \frac{\rho}{1-\rho} \psi_i(\rho) s^i$$
$$= \sum_{j=1}^{\infty} \rho^j \sum_{i=j+2}^{\infty} \psi_{i-j-1}(\rho) s^i + \frac{s}{(1-\rho)(1-\rho s)} + \frac{\rho}{1-\rho} v(\rho, s)$$
$$= \left(\frac{\rho s^2}{1-\rho s} + \frac{\rho}{1-\rho}\right) v(\rho, s) + \frac{s}{(1-\rho)(1-\rho s)},$$
$$v(0, s) = 0 \quad \text{Solving this standard differential equation we obtain}$$

and v(0, s) = 0. Solving this standard differential equation we obtain

$$\nu(\rho, s) = \frac{s(1 - e^{-\rho(s+1)}))}{(s+1)(1-\rho)(1-\rho s)}.$$
(8.12)

We now need to write the Taylor expansion for the latter expression. Using

$$\frac{s}{s+1} = \sum_{m=1}^{\infty} (-1)^{m+1} s^m$$
 and $\frac{1}{1-\rho s} = \sum_{k=0}^{\infty} \rho^k s^k$

yields

$$\frac{s}{s+1}\frac{1}{1-\rho s} = \sum_{l=1}^{\infty} \left(\sum_{k=0}^{l-1} \rho^k (-1)^{l-k+1}\right) s^l = \sum_{l=1}^{\infty} (-1)^{l+1} \left(\sum_{k=0}^{l-1} \rho^k (-1)^{-k}\right) s^l$$
$$= \sum_{l=1}^{\infty} (-1)^{l+1} \frac{1-(-\rho)^l}{1+\rho} s^l = \sum_{l=1}^{\infty} \frac{\rho^l + (-1)^{l+1}}{1+\rho} s^l.$$
(8.13)

Substituting (8.13) and (8.11) into (8.12) gives

$$\begin{split} \nu(\rho,s) \\ &= \frac{s}{s+1} \frac{1}{1-\rho} \frac{1}{1-\rho s} \left(1-e^{-\rho(s+1)}\right) = \frac{1}{1-\rho^2} \left(\sum_{m=1}^{\infty} \left(\rho^m + (-1)^{m+1}\right) s^m \cdot \left(1-e^{-\rho(s+1)}\right)\right) \\ &= \frac{1}{1-\rho^2} \left(\sum_{m=1}^{\infty} \left(\rho^m + (-1)^m\right) s^m - e^{-\rho} \sum_{m=1}^{\infty} \rho^m s^m \sum_{k=0}^{m-1} \frac{(-1)^k}{k!} + e^{-\rho} \sum_{m=1}^{\infty} (-1)^m s^m \sum_{k=0}^{m-1} \frac{\rho^k}{k!}\right), \end{split}$$

which yields (8.10).

Theorem 8.1 provides us with a closed-form but unwieldy expression for the individual throughputs. In case the network size grows to infinity we can obtain a more elegant expression for the throughputs of nodes 1 and 2.

COROLLARY 8.1. As $n \to \infty$,

$$T_1(n) \to 1 - e^{-1}$$
 and $T_2(n) \to e^{-1}$.

8.3 Stability and end-to-end throughput

In this section we relax the assumption that all nodes are saturated and instead consider a multi-hop network where certain buffers may occasionally empty. Specifically, node 1 has an infinite supply of packets available which are forwarded through the network along nodes 2, 3, ..., n. Once transmitted by node n, packets leave the network. Let us denote by $Q_i(t)$ the backlog of node i at time t. Nodes compete for access to the medium as before, with the modification that nodes can only activate when they have packets available for transmission.

We consider a chain of $n = 2\beta + 1$ nodes, $\beta \ge 1$, and we denote by ξ_i the throughput of node *i*. Note that

$$\xi_1 \ge \xi_2 \ge \cdots \ge \xi_n, \tag{8.14}$$

with equality if all nodes are stable. We are interested in the end-to-end throughput ξ_n . In contrast to the saturated case discussed in Section 8.2, the queue lengths play a crucial role, and no explicit expression for the throughputs is known. We shall provide bounds on the end-to-end throughput. To do so we require the following stability results.

LEMMA 8.1. For slotted CSMA with $n = 2\beta + 1$, $\beta \ge 1$, we have

- (i) Nodes $1, 2, \ldots, \beta + 1$ are unstable;
- (ii) Nodes $\beta + 2, \beta + 3, \dots, n$ are stable.

PROOF. Node 1 is saturated by definition, and is thus unstable. Let $i \in \{2, 3, ..., \beta + 1\}$ and assume that node i - 1 is unstable. We will show that

$$\xi_{i-1} > \xi_i, \quad i = 2, 3, \dots, \beta + 1,$$
(8.15)

and conclude by induction that statement (i) holds. In order to demonstrate that (8.15) holds, it is sufficient to show that

$$\mathbb{P}(X_{i-1}(t) = 1) > \mathbb{P}(X_i(t) = 1), \tag{8.16}$$

and that the difference between these two probabilities is bounded away from zero.

To verify (8.16), note that for all $t \ge T$ for some $T < \infty$, node *i* always has a packet for transmission since it is unstable. We let *t* sufficiently large and denote by $A^*(t) = \arg \min_{i:O_i(t)\ge 1} A_i(t)$ the first node to activate in slot *t*, with the convention

8.4 COMPARING SLOTTED AND CONTINUOUS-TIME CSMA

that $A^*(t) = 0$ if all nodes are empty. At most two nodes can be active simultaneously, and by conditioning on $A^*(t)$ we obtain

133

$$\mathbb{P}(X_{i-1}(t) = 1)$$

$$= \mathbb{P}(A^*(t) = i - 1) + \sum_{j=\beta+i}^{2\beta+1} \mathbb{P}(A^*(t) = j) \frac{1}{j - (\beta + i - 1)}$$

$$= \mathbb{P}(A^*(t) = i - 1) + \sum_{j=\beta+i}^{2\beta} \mathbb{P}(A^*(t) = j) \frac{1}{(j + 1) - (\beta + i)} + \mathbb{P}(A^*(t) = 2\beta + 1) \frac{1}{\beta + 2 - i}$$

$$> \left(\mathbb{P}(A^*(t) = i) + \sum_{j=\beta+i+1}^{2\beta+1} \mathbb{P}(A^*(t) = j) \frac{1}{j - (\beta + i)} \right) \mathbb{1}_{\{Q_i(t) \ge 1\}} = \mathbb{P}(X_i(t) = 1),$$

completing the proof of (i).

In order to show statement (ii), let $i \in \{\beta + 2, ..., 2\beta + 1\}$ and assume that node *i* is unstable. Then,

$$\begin{split} \mathbb{P}(X_{i}(t) = 1) &= \mathbb{P}(A^{*}(t) = i) + \sum_{j=1}^{i-\beta-1} \mathbb{P}(A^{*}(t) = j) \frac{1}{\sum_{k=j+\beta+1}^{2\beta+1} \mathbb{1}_{\{Q_{k}(t) \ge 1\}}} \\ &> \left(\mathbb{P}(A^{*}(t) = i-1) + \sum_{j=1}^{i-\beta-2} \mathbb{P}(A^{*}(t) = j) \frac{1}{\sum_{k=j+\beta+1}^{2\beta+1} \mathbb{1}_{\{Q_{k}(t) \ge 1\}}} \right) \\ &= \mathbb{P}(X_{i-1}(t) = 1), \end{split}$$

which is a contradiction, since the throughput of node *i* cannot be greater than that of node i - 1 by (8.14).

Using Lemma 8.1 we can now provide a bound on the end-to-end throughput.

THEOREM 8.2. The end-to-end throughput satisfies $\xi_n > \frac{1}{2\beta+1}$.

The proof of Theorem 8.2 is presented in Appendix 8.A.1. It is based on observing that time may be divided into i.i.d. cycles between instances when nodes $\beta + 2, ..., 2\beta + 1$ empty. The throughput can then be expressed as the ratio of the average number of packets transmitted by node *n* in a typical cycle to the average length of a typical cycle.

8.4 Comparing slotted and continuous-time CSMA

In this section we compare the performance of slotted and continuous-time CSMA. In Section 8.4.1 we inspect the throughput in saturated conditions and in Section 8.4.2 we compare the stability and end-to-end throughput of the multi-hop network.

8.4.1 Saturated networks

Continuous-time and slotted CSMA have been compared in [19], where it is shown that the network-aggregate throughput under the slotted algorithm is lower than for

TIME-SLOTTED CSMA

continuous-time CSMA. The authors then concluded that the CSMA algorithm does not benefit from synchronization. The aggregate throughput is not the only relevant performance measure, however, and we shall now compare fairness under both algorithms.

We consider the saturated CSMA model introduced in Chapter 1, and assume all nodes have equal back-off parameter $v_i = \sigma$, i = 1, 2, ..., n. The network-aggregate throughput in this case is computed in Proposition 5.6. Using this we can plot the aggregate throughput for both continuous-time and slotted CSMA against the number of nodes, see Figure 8.1. This figure shows that for large values of σ , the network-aggregate throughput for slotted CSMA is strictly smaller than for continuous-time CSMA. The reason for this is that slotted CSMA in each slot chooses at random a maximal independent set, while continuous-time CSMA with σ large is typically locked in independent sets of maximum size. Thus continuous-time CSMA on average allows more simultaneous activity, resulting in higher aggregate throughput. This reasoning holds more generally, for example in linear networks with $\beta \ge 2$.



Figure 8.1: The aggregate throughput for continuous-time CSMA (for various values of σ) and slotted CSMA plotted against *n*, for $\beta = 1$.

Thus continuous-time CSMA avoids small maximal independent sets yielding a high aggregate throughput. However, this does not necessarily mean that continuous-time CSMA is better than slotted CSMA: It turns out that the high throughput comes at the cost of unfairness. Figure 8.2 shows the per-node throughputs under saturation for both slotted CSMA (Theorem 8.1) and continuous-time CSMA (Theorem 5.1), for various values of σ . We have seen in this thesis that continuous-time CSMA is unfair, and from Figure 8.2 it is clear that the same holds for slotted CSMA. For small values of σ nodes are active infrequently, and continuous-time CSMA is fairer than slotted CSMA. However, for values of σ that give comparable aggregate throughput to slotted medium access, the continuous-time system is much less fair.

8.4.2 Multi-hop networks

We consider the unsaturated multi-hop network, and compare the end-to-end throughput of slotted CSMA derived in Theorem 8.2 to that of continuous-time CSMA. For

8.5 CONCLUDING REMARKS



135

Figure 8.2: The per-node throughput for continuous-time CSMA (for various values of σ) and slotted CSMA in a network with $\beta = 1$.

continuous-time CSMA we assume all nodes to have back-off rates $\sigma \to \infty$ and we denote by θ_n^* the throughput of node *n* in the multi-hop network. We first show the following stability result for the multi-hop continuous-time CSMA model.

LEMMA 8.2. For continuous-time CSMA with $n = 2\beta + 1$, $\beta \ge 1$,

- (i) Nodes $1, 2, \ldots, \beta + 1$ are unstable;
- (ii) Nodes $\beta + 2, \beta + 3, \dots, 2\beta + 1$ are stable.

The proof of Lemma 8.2 considers the Markov chain embedded at transition instants, and is otherwise analogous to that of Lemma 8.1.

We can now provide upper and lower bounds on the end-to-end throughput for continuous-time CSMA.

THEOREM 8.3. The end-to-end throughput for continuous-time CSMA satisfies

$$\frac{1}{2\beta+2} < \theta_n^* \le \frac{1}{2\beta+1}.$$

The proof of Theorem 8.3 is presented in Appendix 8.A.2.

Combining Theorems 8.2 and 8.3 we see that the end-to-end throughput of slotted CSMA is strictly higher than that of continuous-time CSMA. This is somewhat surprising, in view of the higher aggregate throughput in the continuous-time setting, but can be explained by the better fairness properties of the slotted system.

8.5 Concluding remarks

In this chapter we considered a discrete-time CSMA algorithm, and computed the pernode throughputs. We then studied a multi-hop network and provided a bound on the end-to-end throughput. These results were compared with the performance of continuous-time CSMA, and we observed that for the saturated case continuous-time CSMA has higher aggregate throughput, while slotted CSMA performs better in terms
TIME-SLOTTED CSMA

of fairness. The latter is shown to lead to higher end-to-end throughput for slotted CSMA.

These results suggest an interesting connection between the behavior of the saturated network and the throughput in the unsaturated case. A similar phenomenon was observed in Section 5.5, where it is argued that the back-off rates that provide equal throughputs (for continuous-time CSMA) also perform remarkably well in a multi-hop setting. We conjecture that the minimum and maximum throughput in the saturated case provide lower and upper bounds for the end-to-end throughput in the multi-hop case, respectively. Consequently, in case we have strict fairness (equal throughputs) in the saturated regime, a multi-hop flow that crosses all nodes could attain the saturation throughput.

Appendix

8.A Remaining proofs

8.A.1 Proof of Theorem 8.2

By Lemma 8.1 we have that nodes $\beta + 2, ..., 2\beta + 1$ are stable, so time may be divided into "cycles" that start and end with nodes $\beta + 2, ..., 2\beta + 1$ being empty. Since nodes $1, ..., \beta + 1$ are unstable, the number of packets *R* leaving the network during a cycle and the length of a cycle *U* are identically distributed across cycles. Thus, by renewal reward theory [5] we can express the throughput of our system as the total expected number of packets leaving the system during a cycle divided by the total expected duration of a cycle:

$$\xi_n = \frac{\mathbb{E}[R]}{\mathbb{E}[U]}.$$
(8.17)

A typical cycle will be as follows:

- 1) nodes $1, \ldots, \beta$ jointly finish T_0 transmissions;
- 2) node β + 1 transmits a single packet.

Let τ denote the number of times nodes $\beta + 2, ..., 2\beta + 1$ relinquish access, until all these nodes are empty again, and the cycle ends. Each of these τ events initiates a "sub-cycle" as follows:

for $i = 1, ..., \tau$:

- 3) nodes $L_i, \ldots, \beta + 1$ jointly finish T_i transmissions, N_i of which by $\beta + 1$;
- 4) nodes $\beta + 2, ..., 2\beta + 1$ jointly finish M_i transmissions. At the end of this activity period, node $2\beta + 1$ is empty.

Here L_i is such that $\beta + L_i$ is the rightmost non-empty node at the time when nodes $\beta + 2, ..., 2\beta + 1$ lost access to the channel.

By combining the different components of a cycle we obtain the cycle duration and packet departures as follows

$$R = 1 + N_1 + \dots + N_{\tau} \quad U = T_0 + 1 + (T_1 + M_1) + \dots + (T_{\tau} + M_{\tau}), \tag{8.18}$$

8.A REMAINING PROOFS

Note that nodes $\beta + 2, ..., 2\beta + 1$ experience a joint workload of β for each packet transmitted by node $\beta + 1$, so

$$M_1 + \dots, M_{\tau} = \beta (1 + N_1 + \dots + N_{\tau}).$$
 (8.19)

137

Now, by substituting (8.18) and (8.19) into (8.17) we obtain the following expression for the throughput:

$$\xi_{n} = \frac{\mathbb{E}[1 + N_{1} + \dots, N_{\tau}]}{\mathbb{E}[T_{0} + 1 + (T_{1} + M_{1}) + \dots + (T_{\tau} + M_{\tau})]}$$
$$= \frac{1 + \mathbb{E}[N_{1} + \dots, N_{\tau}]}{2\beta + 1 + \mathbb{E}[T_{1} + \dots + T_{\tau} + \beta(N_{1} + \dots + N_{\tau})]}.$$
(8.20)

The exact state of the system right before the first sub-cycle is known, we have

$$T_0 + 1 \sim \operatorname{geo}\left(\frac{1}{\beta + 1}\right), \quad \mathbb{E}[T_0] = \beta.$$
 (8.21)

We know that $L_1 = 2$ and $2 \le L_i \le \beta$, $i = 1, ..., \tau$. We say there are a total of H_i non-empty nodes to the right from $\beta + 1$, $1 \le H_i \le \beta - 1$. So only nodes $L_i, ..., \beta + 1$ can win the next competition without a node on the right gaining access. From this we know

$$T_i + 1 \sim \text{geo}(\frac{H_i + L_i - 1}{\beta - L_i + 2}), \quad N_i + 1 \sim \text{geo}(\frac{1}{\beta - L_i + 2}).$$

Using this, it can be seen that

$$\mathbb{E}[T_i] = (\beta - L_i + 2)\mathbb{E}[N_i] = \frac{\beta + L_i + 2}{H_i + L_i - 1}.$$
(8.22)

Equation (8.22) implies, in particular, that

$$(\beta + 1)\mathbb{E}[N_i] - \mathbb{E}[T_i] = \frac{L_i - 1}{M_i + L_i - 1} \le 1.$$
(8.23)

From stability of nodes β + 2, . . . , 2 β + 1 we know τ < ∞ . Moreover, it holds that

$$\tau \le 1 + N_1 + \ldots + N_{\tau}. \tag{8.24}$$

This is true because a sub-cycle always ends by a successful transmission of node $2\beta + 1$. This implies that the number of sub-cycles may not be larger than the number of packets leaving the system during the entire cycle.

Substituting this into (8.20) yields the lower bound

$$\begin{split} \xi_n &= \frac{1 + \mathbb{E}[N_1 + \ldots + N_{\tau}]}{2\beta + 1 + \mathbb{E}[T_1 + \ldots + T_{\tau} + \beta(N_1 + \ldots + N_{\tau})]} \\ &= \frac{1 + \mathbb{E}[N_1 + \ldots + N_{\tau}]}{2\beta + 1 + \mathbb{E}[(\beta - L_1 + 2)N_1 + \ldots + (\beta - L_{\tau} + 2)N_{\tau}] + \beta\mathbb{E}[N_1 + \ldots + N_{\tau}]} \\ &\geq \frac{1 + \mathbb{E}[N_1 + \ldots + N_{\tau}]}{2\beta + 1 + (2\beta + 1)\mathbb{E}[N_1 + \ldots + N_{\tau}]} = \frac{1}{2\beta + 1}, \end{split}$$

completing the proof.

TIME-SLOTTED CSMA

8.A.2 Proof of Theorem 8.3

A proof of this theorem may be given following the steps of the proof of Theorem 8.2. Indeed, similar to the slotted case, we see that nodes $\beta + 2, ..., 2\beta + 1$ are stable and nodes $1, ..., \beta + 1$ are unstable, by Lemma 8.2. Thus, the time may be divided into "cycles" that start and end with nodes $\beta + 2, ..., 2\beta + 1$ being empty, and again the throughput of our system is equal to the total expected number of packets leaving the system during a cycle divided by the total expected duration of a cycle.

A typical cycle is constructed in exactly the same way as in the time-slotted case, with the addition that a cycle is extended with a residual transmission time of node $2\beta + 1$. Now, instead of (8.20) in the proof of Theorem 8.2, we have

$$\theta_n^* = \frac{1 + \mathbb{E}[N_1 + \dots, N_{\tau}]}{2\beta + 1 + \mathbb{E}[T_1 + \dots + T_{\tau} + \beta(N_1 + \dots + N_{\tau}) + \tau]}.$$
(8.25)

Note that in the continuous-time case the lengths of the various parts of a typical cycle are no longer geometrically distributed, but their expectations are exactly the same as in the time-slotted case. Hence, (8.22)-(8.24) still hold. Taking this into account, we substitute (8.22) and (8.24) into (8.25) to obtain a lower bound:

$$\begin{split} \theta_n^* &= \frac{1 + \mathbb{E}[N_1 + \ldots + N_{\tau}]}{2\beta + 1 + \mathbb{E}[T_1 + \ldots + T_{\tau} + \beta(N_1 + \ldots + N_{\tau}) + \tau]} \\ &\geq \frac{1 + \mathbb{E}[N_1 + \ldots + N_{\tau}]}{2\beta + 2 + (2\beta + 1)\mathbb{E}[N_1 + \ldots + N_{\tau}]} \\ &\geq \lim_{y \to 0} \frac{1 + y}{2\beta + 2 + (2\beta + 1)y} = \frac{1}{2\beta + 2}, \end{split}$$

By (8.23) we have that

$$\begin{split} \theta_n^* &= \frac{1 + \mathbb{E}[N_1 + \ldots + N_{\tau}]}{2\beta + 1 + \mathbb{E}[T_1 + \ldots + T_{\tau} + \beta(N_1 + \ldots + N_{\tau}) + \tau]} \\ &= \frac{1 + \mathbb{E}[N_1 + \ldots + N_{\tau}]}{2\beta + 1 + \mathbb{E}[T_1 - (\beta + 1)N_1 + \ldots + T_{\tau} - (\beta + 1)N_{\tau} + (2\beta + 1)(N_1 + \ldots + N_{\tau}) + \tau]} \\ &\leq \frac{1 + \mathbb{E}[N_1 + \ldots + N_{\tau}]}{(2\beta + 1)(1 + \mathbb{E}[N_1 + \ldots + N_{\tau}])} = \frac{1}{2\beta + 1}, \end{split}$$

which is the upper bound.

BIBLIOGRAPHY

- 802.11. IEEE standards for information technology Telecommunications and information exchange between systems - local and metropolitan area network
 specific requirements - Part 11: wireless LAN medium access control (MAC) and physical layer (PHY) specifications, 1999.
- [2] N. Abramson. The ALOHA system another alternative for computer communications. In *Proc. AFIPS*, pages 281–285, 1970.
- [3] V. Anantharam. The stability region of the finite-user slotted ALOHA protocol. *IEEE Trans. Inf. Theory*, 37(3):535–540, 1991.
- [4] D.M. Andrews, K. Kumaran, K. Ramanan, A.L. Stolyar, R. Vijayakumar, and P. Whiting. Scheduling in a queueing system with asynchronously varying service rates. *Prob. Eng. Inf. Sc.*, 18(2):191–217, 2004.
- [5] S. Asmussen. *Applied Probability and Queues*. Springer-Verlag, New York, second edition, 2003.
- [6] Y Baryshnikov, E.G. Coffman, Jr., and P.R. Jelenković. Space filling and depletion. *J. Appl. Probab.*, 41(3):691–702, 2004.
- [7] G. Bianchi. Performance analysis of the IEEE 802.11 distributed coordination function. *IEEE J. Sel. Areas Commun.*, 18(3):535–547, 2000.
- [8] T. Bonald. The Erlang model with non-Poisson call arrivals. In *Proc. ACM SIG-METRICS*, pages 276–286, Saint-Malo, France, June 26–30 2006.
- [9] T. Bonald and M. Feuillet. On the stability of flow-aware CSMA. *Performance Evaluation*, 67(11):1219–1229, 2010.
- [10] T. Bonald and M. Feuillet. On flow-aware CSMA in multi-channel wireless networks. In *Proc. CISS*, Baltimore, MD, March 23–25 2011.
- [11] R.R. Boorstyn and A. Kershenbaum. Throughput analysis of multihop packet radio. In *Proc. ICC*, pages 1361–1366, 1980.
- [12] R.R. Boorstyn, A. Kershenbaum, B. Maglaris, and V. Sahin. Throughput analysis in multihop CSMA packet radio networks. *IEEE Trans. Commun.*, 35(3):267–274, 1987.

BIBLIOGRAPHY

- [13] S.C. Borst, O.J. Boxma, and P.R. Jelenković. Reduced-load equivalence and induced burstiness in GPS queues with long-tailed traffic flows. *Queueing Syst.*, 43(4):273–306, 2003.
- [14] E. Brockmeyer, H.L. Halstrom, and A. Jensen. *The life and works of A. K. Erlang.* Transactions of the Danish Academy of Technical Sciences, Copenhagen, first edition, 1948.
- [15] P. Chaporkar and S. Sarkar. Stable scheduling policies for maximizing throughput in generalized constrained queueing systems. *IEEE Trans. Automat. Control*, 53(8):1913–1931, 2008.
- [16] N.G. de Bruijn. *Asymptotic methods in analysis*. Dover Publications Inc., New York, third edition, 1981.
- [17] D. Denteneer, S.C. Borst, P.M. van de Ven, and G. Hiertz. IEEE 802.11s and the philosophers' problem. *Statistica Neerlandica*, 62(3):283–298, 2008.
- [18] O. Dousse. Revising buffering in multihop CSMA/CA wireless networks. In *Proc. Secon*, pages 580–589, San Diego, CA, June 18–21 2007.
- [19] M. Durvy and P. Thiran. A packing approach to compare slotted and non-slotted medium access control. In *Proc. Infocom*, Barcelona, Spain, April 23–29 2006.
- [20] M. Durvy, O. Dousse, and P. Thiran. Modeling the 802.11 protocol under different capture and sensing capabilities. In *Proc. Infocom*, pages 2356–2360, Anchorage, AK, May 6–12 2007.
- [21] M. Durvy, O. Dousse, and P. Thiran. Border effects, fairness, and phase transitions in large wireless networks. In *Proc. Infocom*, pages 601–609, Phoenix, AZ, April 13–18 2008.
- [22] M. Durvy, O. Dousse, and P. Thiran. Self-organization properties of CSMA/CA systems and their consequences on fairness. *IEEE Trans. Inf. Theory*, 55(3): 931–943, 2009.
- [23] A. Eryilmaz and R. Srikant. Fair resource allocation in wireless networks using queue-length-based scheduling and congestion control. In *Proc. Infocom*, pages 1794–1803, Miami, FL, March 13–17 2005.
- [24] A. Eryilmaz, R. Srikant, and J.R. Perkins. Stable scheduling policies for fading wireless channels. *IEEE/ACM Trans. Netw.*, 13(2):411–424, 2005.
- [25] C. Fricker and M.R. Jaibi. Monotonicity and stability of periodic polling models. *Queueing Syst.*, 15(1–4):211–238, 1994.
- [26] L. Fu, S.C. Liew, and J. Huang. Effective carrier sensing in CSMA networks under cumulative interference. In *Proc. Infocom*, San Diego, CA, March 15–19 2010.
- [27] M. Garetto, T. Salonidis, and E.W. Knightly. Modeling per-flow throughput and capturing starvation in CSMA multi-hop wireless networks. *IEEE/ACM Trans. Netw.*, 16(4):864–877, 2008.

- [28] L. Georgiadis, M.J. Neely, and L. Tassiulas. Resource allocation and cross-layer control in wireless networks. *Found. Trends Netw.*, 1(1):1–144, 2006.
- [29] J. Ghaderi and R. Srikant. On the design of efficient CSMA algorithms for wireless networks. In *Proc. CDC*, pages 954–959, Atlanta, GA, December 15–17 2010.
- [30] P. Gupta and P.R. Kumar. The capacity of wireless networks. *IEEE Trans. Inf. Theory*, 46(2):388–404, 2000.
- [31] T. Hellings, J.S.H. van Leeuwaarden, S.C. Borst, and D. Denteneer. Extra backoff flow control in multi-hop wireless networks. *Performance Evaluation*, 68(8): 728-746, 2011.
- [32] L. Jiang and J. Walrand. A distributed CSMA algorithm for throughput and utility maximization in wireless networks. In *Proc. 46th Annual Allerton Conf. Commun., Control and Comp.*, pages 1511–1519, Monticello, IL., September 23– 26 2008.
- [33] L. Jiang and J. Walrand. Convergence and stability of a distributed CSMA algorithm for maximal network throughput. Technical Report http://www. eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-43.pdf, UC Berkeley, 2009.
- [34] L. Jiang and J. Walrand. A distributed CSMA algorithm for throughput and utility maximization in wireless networks. *IEEE/ACM Trans. Netw.*, 18(3):960– 972, 2010.
- [35] L. Jiang and J. Walrand. Approaching throughput-optimality in distributed CSMA scheduling algorithms with collisions. *IEEE/ACM Trans. Netw. (to appear)*, 2011.
- [36] L. Jiang, D. Shah, J. Shin, and J. Walrand. Distributed random access algorithm: scheduling and congestion control. *IEEE Trans. Inf. Theory*, 56(12):6182–6207, 2010.
- [37] N. Kahale and P.E. Wright. Dynamic global packet routing in wireless networks. In *Proc. Infocom*, pages 1416–1423, Kobe, Japan, April 7–12 1997.
- [38] F.P. Kelly. Reversibility and Stochastic Networks. Wiley, Chichester York, 1979.
- [39] F.P. Kelly. Stochastic models of computer communication systems. *J. Roy. Stat. Soc. B*, 47(3):379–395, 1985.
- [40] F.P. Kelly. One-dimensional circuit-switched networks. *Ann. Probab.*, 15(3): 1166-1179, 1987.
- [41] F.P. Kelly. Loss networks. Ann. Appl. Probab., 1(3):319-378, 1991.
- [42] A. Kershenbaum, R.R. Boorstyn, and M. Chen. An algorithm for evaluation of throughput in multihop packet radio. *IEEE J. Sel. Areas Commun.*, 5(6):1003–1012, 1987.

141

BIBLIOGRAPHY

- [43] T.H. Kim, J. Ni, R. Srikant, and N.H. Vaidya. On the achievable throughput of CSMA under imperfect carrier sensing. In *Proc. Infocom*, pages 1674–1682, Shanghai, China, April 10–15 2011.
- [44] L. Kleinrock and F.A. Tobagi. Packet switching in radio channels: part I carrier sense multiple-access modes and their throughput-delay characteristics. *IEEE Trans. Commun.*, 23(12):1400–1416, 1975.
- [45] J. Lee, J. Lee, Y. Yi, S. Chong, A. Proutière, and M. Chiang. Implementing utilityoptimal CSMA. In *Proc. 47th Annual Allerton Conf. Commun., Control and Comp.*, pages 102–111, Monticello, IL, September 30–October 2 2009.
- [46] M. Lelarge. Asymptotic behavior of generalized processor sharing queues under subexponential assumptions. *Queueing Syst.*, 62(1–2):51–73, 2009.
- [47] D.A. Levin, Y. Peres, and E.L. Wilmer. *Markov chains and mixing times*. American Mathematical Society, first edition, 2008.
- [48] S-Y.R. Li, R.W. Yeung, and C. Ning. Linear network coding. *IEEE Trans. Inf. Theory*, 49(2):371–381, 2003.
- [49] S.C. Liew, C. Kai, J. Leung, and B. Wong. Back-of-the-envelope computation of throughput distributions in CSMA wireless networks. *IEEE Trans. Mobile Comput.*, 9(9):1319–1331, 2010.
- [50] T.-Y. Lin and J.C. Hou. Interplay of spatial reuse and SINR-determined data rates in CSMA/CA-based, multi-hop, multi-rate wireless networks. In *Proc. Infocom*, pages 803–811, Anchorage, AK, May 6–12 2007.
- [51] X. Lin and N.B. Shroff. The impact of imperfect scheduling on cross-layer rate control in wireless networks. In *Proc. Infocom*, pages 1804–1814, Miami, FL, March 13–17 2005.
- [52] X. Lin, N.B. Shroff, and R. Srikant. On the connection-level stability of congestion-controlled communication networks. *IEEE Trans. Inf. Theory*, 54 (5):2317–2338, 2008.
- [53] J. Liu, Y. Yi, A. Proutière, M. Chiang, and H.V. Poor. Towards utility-optimal random access without message passing. *Wirel. Commun. Mob. Comput.*, 10(1): 115–128, 2010.
- [54] S. Liu, L. Ying, and R. Srikant. Scheduling in multichannel wireless networks with flow-level dynamics. In *Proc. ACM SIGMETRICS*, pages 191–202, New York, NY, June 14–18 2010.
- [55] S. Liu, L. Ying, and R. Srikant. Throughput-optimal opportunistic scheduling in the presence of flow-level dynamics. *IEEE/ACM Trans. Netw.*, 19(4):1057–1070, 2011.
- [56] M. Lotfinezhad and P. Marbach. Throughput-optimal random access with orderoptimal delay. In *Proc. Infocom*, pages 2867–2875, Shanghai, China, April 10–15 2011.

- [57] H. Ma, R. Vijaykumar, S. Roy, and J. Zhu. Optimizing 802.11 wireless mesh networks based on physical carrier sensing. *IEEE/ACM Trans. Netw.*, 17(5):1550–1563, 2009.
- [58] N. McKeown, V. Anantharam, and J. Walrand. Achieving 100% throughput in an input-queued switch. In *Proc. Infocom*, pages 296–302, San Francisco, CA, March 24–28 1996.
- [59] N. McKeown, A. Mekkittikul, V. Anantharam, and J. Walrand. Achieving 100% throughput in an input-queued switch. *IEEE Trans. Commun.*, 47(8):1260–1267, 1999.
- [60] K. Medepalli and F.A. Tobagi. Towards performance modeling of IEEE 802.11 based wireless networks: A unified framework and its applications. In *Proc. Infocom*, Barcelona, Spain, April 23–26 2006.
- [61] S.P. Meyn. *Control Techniques for Complex Networks*. Cambridge University Press, 2007.
- [62] S.P. Meyn and R.L. Tweedie. *Markov Chains and Stochastic Stability*. Springer Verlag, first edition, 1993.
- [63] C. Moallemi and D. Shah. On the flow-level dynamics of a packet-switched network. In *Proc. ACM SIGMETRICS*, pages 83–94, New York, NY, June 14–18 2010.
- [64] A. Müller and D. Stoyan. *Comparison Methods for Stochastic Models and Risks*. Wiley, 2002.
- [65] B. Nardelli, J. Lee, K. Lee, Y. Yi, S. Chong, E.W. Knightly, and M. Chiang. Experimental evaluation of optimal CSMA. In *Proc. Infocom*, pages 1188–1196, Shanghai, China, April 10–15 2011.
- [66] M.J. Neely. Energy optimal control for time-varying wireless networks. *IEEE Trans. Inf. Theory*, 52(7):2915–2934, 2006.
- [67] M.J. Neely, E. Modiano, and C.E. Rohrs. Dynamic power allocation and routing for time-varying wireless networks. *IEEE J. Sel. Areas Commun.*, 23(1):89–103, 2005.
- [68] M.J. Neely, E. Modiano, and C.-P. Li. Fairness and optimal stochastic control for heterogeneous networks. *IEEE/ACM Trans. Netw.*, 16(2):396–409, 2008.
- [69] M.F. Neuts. *Matrix-Geometric Solutions in Stochastic Models: an Algorithmic Approach*. Dover Publications Inc., 1981.
- [70] J. Ni, B. Tan, and R. Srikant. Q-CSMA: queue-length based CSMA/CA algorithms for achieving maximum throughput and low delay in wireless networks. In *Proc. Infocom*, San Diego, CA, March 15–19 2010.
- [71] E. Pinsky and Y. Yemini. The asymptotic analysis of some packet radio networks. *IEEE J. Sel. Areas Commun.*, 4(6):938–945, 1986.



BIBLIOGRAPHY

- [72] S. Rajagopalan, D. Shah, and J. Shin. Network adiabatic theorem: an efficient randomized protocol for content resolution. In *Proc. ACM SIGMET-RICS/Performance*, pages 133–144, Seattle, WA, June 15–19 2009.
- [73] L.G. Roberts. ALOHA packet system with and without slots and capture. *ACM SIGCOMM Comput. Commun. Rev.*, 5(2):28–42, 1975.
- [74] B. Sadiq and G. de Veciana. Throughput optimality of delay-driven maxweight scheduler for a wireless system with flow dynamics. In *Proc. 47th Annual Allerton Conf. Commun., Control, Comp.*, pages 1097–1102, Monticello, IL., September 30–October 2 2009.
- [75] D. Shah and J. Shin. Delay optimal queue-based CSMA. http://www.cc. gatech.edu/~jshin72/csma_delay_ss.pdf (preprint), 2010.
- [76] D. Shah and J. Shin. Efficient queue-based CSMA with collisions. http://arxiv. org/abs/arXiv:1003.2749 (preprint), 2010.
- [77] G. Sharma, R.R. Mazumdar, and N.B. Shroff. On the complexity of scheduling in wireless networks. In *Proc. MobiCom*, pages 227–238, Los Angeles, CA, September 24–29 2006.
- [78] G. Sharma, N.B. Shroff, and R.R. Mazumdar. Joint congestion control and distributed scheduling for throughput guarantees in wireless networks. In *Proc. Infocom*, pages 2072–2080, Anchorage, AK, May 6–12 2007.
- [79] J. Shi, O. Gurewitz, V. Manusco, J. Camp, and E.W. Knightly. Measurement and modeling of the origins of starvation in congestion controlled mesh networks. In *Proc. Infocom*, pages 1633–1641, Phoenix, AZ, April 13–18 2008.
- [80] V. Shneer and P.M. van de Ven. Comparing slotted and continuous CSMA: throughputs and fairness. In *Proc. Performance*, Amsterdam, NL, October 18–20 2011.
- [81] A.L. Stolyar. MaxWeight scheduling in a generalized switch: state space collapse and workload minimization in heavy traffic. *Ann. Appl. Probab.*, 14(1):1–53, 2004.
- [82] A.L. Stolyar. Maximizing queueing network utility subject to stability: greedy primal dual algorithm. *Queueing Syst.*, 50(4):401–457, 2005.
- [83] A.L. Stolyar. Greedy primal-dual algorithm for dynamic resource allocation in complex networks. *Queueing Syst.*, 54(3):203–220, 2006.
- [84] Y. Suhov and U.A. Rozikov. A hard-core model on a Cayley tree: An example of a loss network. *Queueing Syst.*, 46(1–2):197–212, 2004.
- [85] W. Szpankowski. Stability conditions for multidimensional queueing systems with computer applications. *Oper. Res.*, 36(6):944–957, 1988.
- [86] W. Szpankowski. Stability conditions for some distributed systems: Buffered random access systems. *Adv. Appl. Probab.*, 26(2):498–515, 1994.

- [87] L. Tassiulas. Linear complexity algorithms for maximum throughput in radio networks and input queued switches. In *Proc. Infocom*, pages 533–539, San Francisco, CA, March 29–April 2 1998.
- [88] L. Tassiulas and A. Ephremides. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Trans. Automat. Contr.*, 37(12):1936–1948, 1992.
- [89] L. Tassiulas and A. Ephremides. Dynamic server allocation to parallel queues with randomly varying connectivity. *IEEE Trans. Inf. Theory*, 39(2):466–478, 1993.
- [90] P.M. van de Ven, S.C. Borst, and V. Shneer. Instability of MaxWeight scheduling algorithms. In *Proc. Infocom*, pages 1701–1709, Rio de Janeiro, Brazil, April 19–25 2009.
- [91] P.M. van de Ven, J.S.H. van Leeuwaarden, D. Denteneer, and A.J.E.M. Janssen. Spatial fairness in wireless multi-access networks. In *Proc. ValueTools*, Pisa, Italy, October 20–23 2009.
- [92] P.M. van de Ven, S.C. Borst, J.S.H. van Leeuwaarden, and A. Proutière. Insensitivity and stability of random-access networks. *Performance Evaluation*, 67(11): 1230–1242, 2010.
- [93] P.M. van de Ven, A.J.E.M. Janssen, and J.S.H. van Leeuwaarden. Optimal tradeoff between exposed and hidden nodes in large wireless networks. In *Proc. ACM SIGMETRICS*, pages 179–190, New York, NY, June 14–18 2010.
- [94] P.M. van de Ven, S.C. Borst, and L. Ying. Spatial inefficiency of MaxWeight scheduling. In *Proc. WiOpt*, pages 62–69, Princeton, NJ, May 9–13 2011.
- [95] P.M. van de Ven, A.J.E.M. Janssen, J.S.H. van Leeuwaarden, and S.C. Borst. Achieving target throughputs in random-access networks. *Performance Evaluation (to appear)*, 2011.
- [96] P.M. van de Ven, J.S.H. van Leeuwaarden, D. Denteneer, and A.J.E.M. Janssen. Spatial fairness in linear random-access networks. *Performance Evaluation (to appear)*, 2011.
- [97] N.M. van Dijk. *Queueing networks and product forms : a systems approach*. Wiley, Chichester, 1993.
- [98] X. Wang and K. Kar. Throughput modelling and fairness issues in CSMA/CA based ad-hoc networks. In *Proc. Infocom*, pages 23–34, Miami, FL, March 13–17 2005.
- [99] F.F. Wu and C.A. Desoer. Global inverse function theorem. *IEEE Trans. Circuit Theory*, 19(2):199–201, 1972.
- [100] X. Wu and R. Srikant. Regulated maximal matching: a distributed scheduling algorithm for multi-hop wireless networks with node-exclusive spectrum sharing. In *Proc. 44th IEEE CDC-ECC*, pages 5342–5347, Seville, Spain, December 12–15 2005.



- [101] X. Wu and R. Srikant. Scheduling efficiency of distributed greedy scheduling algorithms in wireless networks. *IEEE Trans. Mobile Comput.*, 6(6):595–605, 2007.
- [102] S. Xu and T. Saadawi. Does the IEEE 802.11 MAC protocol work well in multihop wireless ad hoc networks? *IEEE Commun. Mag.*, 39(6):130–137, 2001.
- [103] F. Xue and P.R. Kumar. Scaling laws for ad-hoc wireless networks: An information theoretic approach. *Found. Trends Netw.*, 1(2):145–270, 2006.
- [104] X. Yang and N. Vaidya. On physical carrier sensing in wireless ad hoc networks. In *Proc. Infocom*, pages 2525–2535, Miami, FL, March 13–17 2005.
- [105] S. Zachary and I. Ziedins. Loss networks and Markov random fields. *J. Appl. Probab.*, 36(2):403–414, 1999.
- [106] M. Zafer and E. Modiano. Blocking probability and channel assignment in wireless networks. *IEEE Trans. Wireless Commun.*, 5(4):869–879, 2006.
- [107] H. Zhai and Y. Fang. Physical carrier sensing and spatial reuse in multirate and multihop wireless ad hoc networks. In *Proc. Infocom*, Barcelona, Spain, April 23–29 2006.

SUMMARY

Next-generation wireless networks will likely evolve from cellular and small-scale home networks to large, inter-connected networks that form the backbone for lowcost internet access. A defining characteristic of wireless networks is that all nodes share the same medium for their transmissions, and consequently, simultaneous transmissions from nearby nodes will interfere with each other. The resulting performance issues can be mitigated by regulating node activity.

Various mechanisms exist for regulating node access to the wireless medium, which can be categorized into *scheduled-access* and *random-access* algorithms. The former involve a centralized entity that controls the behavior of all nodes, while random-access constitutes a class of randomized, distributed algorithms. Random-access algorithms are popular for their simplicity, and their distributed nature makes them well-suited for large, dynamic wireless networks. Scheduled-access networks generally have better performance since the presence of an omniscient controller allows for coordination between nodes, but entail higher implementation complexity.

One well-known algorithm for centralized access is MaxWeight scheduling, which is popular for its ability to achieve maximum stability and throughput optimality in a wide variety of scenarios. The distinguishing characteristic of MaxWeight policies is that these require solving the maximum weighted independent set problem of the underlying interference graph. The maximum-stability guarantees however rely on the premise that the system consists of a fixed set of flows, while in reality the collection of active flows dynamically varies. In Chapters 2 and 3 we demonstrate that in the presence of flow-level dynamics the algorithm may no longer be throughput-optimal, and we identify two causes for the instability: (i) failure to fully exploit rate variations; and (ii) spatial inefficiency.

In Chapter 2 we consider the MaxWeight scheduling algorithm in a single downlink scenario with varying transmission rates. We identify a simple necessary and sufficient condition for stability, and show that MaxWeight policies may fail to provide maximum stability. The intuitive explanation is that these policies tend to favor flows with large backlogs, so that the rate variations of flows with smaller backlogs are not fully exploited.

The second cause for instability is studied in Chapter 3, where we consider a spatial setting in which flows arrive at random in some finite space, and multiple flows may be scheduled simultaneously, subject to certain interference constraints. The MaxWeight scheduler tends to serve flows with large backlogs, even when the resulting spatial reuse is not particularly efficient. We show that the inability of MaxWeight policies to exploit maximum spatial reuse patterns may lead to instability.

SUMMARY

Random-access algorithms were originally designed for symmetric deployment scenarios, under the assumption that all nodes interfere with each other. The recent trend towards large-scale distributed networks has induced a shift from such full interference scenarios to networks with partial interference graphs, which has brought to light many performance issues that require more study and give rise to major mathematical challenges. A particularly popular random-access algorithm is carrier-sense multiple-access (CSMA), which is implemented for example in the widely deployed IEEE 802.11 standard. This protocol reduces interference by introducing a carrier-sensing mechanism that allows nodes to transmit only when nearby nodes are inactive. In Chapters 4-8 we study the CSMA algorithm.

In recent years relatively parsimonious models have emerged that provide a useful tool in evaluating the throughput characteristics of CSMA-like networks. These models essentially assume that the interference constraints can be represented by a general conflict graph, and that the various nodes activate asynchronously after an exponential back-off time whenever none of their neighbors are active. It turns out that the assumption of exponential transmission times and back-off durations can be relaxed, as we show in Chapter 4. We also consider the unsaturated model, where buffers may occasionally be empty as packets are randomly generated and transmitted over time. We explicitly identify the stability conditions for the complete interference graph, and illustrate the difficulties that arise for partial interference graphs.

In Chapters 5-7 we study the throughput of random-access networks using the CSMA model. Such networks may exhibit severe unfairness, in the sense that some nodes receive consistently higher throughput than others. In Chapter 5 we study this phenomenon in linear networks, and remove the unfairness completely by choosing node-specific mean back-off times. We obtain explicit expressions for the fair back-off times and the resulting throughput.

The more general problem of finding the mean back-off times that yield a certain throughput vector is addressed in Chapter 6. In order to compute the required backoff times, we show that the throughput function is globally invertible, and we present several numerical procedures for calculating this inverse, based on fixed-point iteration and Newton's method. The ability to determine the network parameters that yield a certain throughput vector allows for much more flexible design of wireless networks.

The carrier-sensing mechanism of CSMA blocks all nodes within a certain sensing range of an active node from transmitting. This mechanism reduces collisions, but also introduces a complex tradeoff for the choice of the sensing range. When the sensing range increases, the interference is reduced, but so is spatial reuse. In Chapter 7 we study this tradeoff in a linear network, and determine the throughput-optimal sensing range. We show that the value of the optimal sensing range depends on the mean back-off times of the nodes.

Finally, in Chapter 8 we consider a time-slotted version of the CSMA algorithm in a linear network. We compute the aggregate throughput and per-node throughputs under saturation conditions, as well as stability and end-to-end throughput for an unsaturated multi-hop network. These results are compared to those obtained for the continuous-time CSMA model.

ABOUT THE AUTHOR

Peter van de Ven was born in Oss, The Netherlands, on July 22, 1984. He completed secondary education at Maaslandcollege, Oss, in 2002. In September 2002 he started his studies in Applied Mathematics at Eindhoven University of Technology, The Netherlands, where he obtained his B.Sc. degree in 2005 and his M.Sc. degree in 2007. His master's thesis on wireless mesh networks was written during a nine-month internship at Philips Research, Eindhoven. Afterwards, he started working towards his Ph.D. at Eindhoven University of Technology and Eurandom. Under the supervision of Sem Borst and Johan van Leeuwaarden he worked on stochastic models for wireless networks, and this research led to various publications in international conferences and journals, and this thesis. During the winter of 2010-2011 he spent three months at Technicolor Paris Research Lab. Peter defends his thesis on December 19, 2011. As of October 2011 he is the Goldstine Postdoctoral Fellow at the Business Analytics and Mathematical Sciences Department of the IBM Thomas J. Watson Research Center, Yorktown Heights, U.S.