LEIBNIZ UNIVERSITÄT HANNOVER
FAKULTÄT FÜR ELEKTROTECHNIK UND INFORMATIK

A Measurement-Based Approach to Service Modeling and Bandwidth Estimation in IEEE 802.11 Wireless Networks

Von der Fakultät für Elektrotechnik und Informatik
der Gottfried Wilhelm Leibniz Universität Hannover
zur Erlangung des akademischen Grades

Doktor-Ingenieur

genehmigte

Dissertation

von

Dipl.-Ing. Michael Bredel
geboren am 07. Januar 1979 in Bingen

2012

Referent	:	Prof. Dr.-Ing. Markus Fidler
Korreferent	:	Prof. Dr.-Ing. Ralf Steinmetz
Tag der Promotion	:	24. Februar 2012

ABSTRACT

This thesis deals with the performance analysis of IEEE 802.11 wireless networks. Based on an analysis of the distributed coordination function, which aims at an efficient medium access, the thesis addresses fairness issues that occur in the case of several wireless stations contending for a channel. In light of the success of IEEE 802.11, it is remarkable that there is little consensus on the actual degree of fairness achieved, particularly bearing in mind its impact on quality of service. Moreover, we evaluate service estimation methods for wireless networks. Based on our findings and in combination with our fairness model, we show how a station can estimate its fair bandwidth share from passive measurements of traffic arrivals and departures.

In Chapters 4 to 5 of this thesis we provide an accurate model for the fairness of the DCF. Given M greedy stations, we assume fairness if a tagged station contributes a share of $1/M$ to the overall number of transmitted packets. We derive the probability distribution of fairness deviations and support our analytical results with an extensive set of measurements. We find a closed-form expression for the improvement of long-term over short-term fairness. Regarding the random countdown values, we quantify the significance of their distribution whereas we discover that fairness is largely insensitive to the distribution parameters. Based on our findings, we view the DCF as emulating an ideal fair queuing system to quantify the deviations from a fair rate allocation. We deduce a stochastic service curve model for the DCF to predict packet delays in IEEE 802.11 that can have a tremendous impact on quality of service and measurements.

These aspects are further explored in Chapter 6. Here we present results from an extensive measurement study of wireless bandwidth estimation in IEEE 802.11 WLANs using the DCF. We show that a number of known iterative probing methods, which are based on the assumption of first-come

first-serve scheduling, can be expected to report the fair bandwidth share of a new flow rather than the available bandwidth. Our measurement results confirm this view and we conclude that under the current probe gap and probe rate models the fair share can only be loosely related to the available bandwidth. Like several other studies we report that packet sizes have a tremendous impact on bandwidth estimates. Unlike these studies we can, however, show that minor modifications to known methods for wired networks, such as Pathload, can solve previously indicated limitations of these methods in wireless networks. Moreover, we develop and analyze a new bandwidth estimation method for wireless networks. To this end, we employ active probing and continuously inject packet probes into the network in order to detect changes in the bandwidth process over time. We present the key challenges and analyze the trade-offs between fast change detection and estimate smoothness. We show the benefit of using Kalman filtering to obtain optimal results under certain conditions and provide a way to parameterize the filter with respect to specific use cases. Furthermore, we evaluate the influence of probing-train length on the results. Finally, we show how user space applications, such as video streaming, can benefit from these measurements.

ZUSAMMENFASSUNG

Die vorliegende Dissertation befasst sich mit der Leistungsanalyse draht-
loser Netzwerke auf Basis des IEEE 802.11 Standards. Basierend auf einer
Analyse der Distributed Coordination Function, die einen dezentralen Me-
dienzugriff erlaubt, werden Fairness-Aspekte untersucht die beim Zusam-
menspiel mehrerer drahtloser Stationen die Leistungsfähigkeit des Net-
zes maßgeblich beeinflussen. Angesichts der weiten Verbreitung ist es
bemerkenswert, dass bis zum heutigen Zeitpunkt keine Einigkeit über
den Grad der erreichten Fairness besteht. Im Verlauf dieser Arbeit werden
zunächst bestehende Methoden zum Messen von Dienstgüteparametern in
drahtlosen Netzen evaluiert. Auf der Grundlage dieser Erkenntnisse und
in Kombination mit einem zuvor entwickeltem Fairness-Modell wird an-
schließend eine neue Methode zum Abschätzen der verfügbaren Bandbreite
vorgestellt.

In den Kapiteln 4 und 5 wird ein Modell entwickelt, welches die Fair-
ness der DCF beschreibt. Dabei wird ein Netzwerk mit M kontinuierlich
sendenden Stationen als perfekt fair angenommen, wenn eine Station einen
Anteil von $1/M$ aller gesendeten Pakete übertragen kann. Ausgehend von
dieser Fairnessdefinition wird die Wahrscheinlichkeitsverteilung bezüglich
Abweichungen von dieser Fairness modelliert. Es werden Formeln in
geschlossener Form angegeben, die die Fairness-Verbesserung von langen
Beobachtungszeiträume gegenüber kurzen Zeiträumen beschreiben. Die
Ergebnisse werden durch Messungen untermauert und bestätigt. Bezüglich
der in der DCF verwendeten Backoff-Werte wird die Bedeutung der zuge-
hörigen Wahrscheinlichkeitsverteilung und deren Einfluss auf die erzielte
Fairness herausgearbeitet. Basierend auf den gewonnenen Erkenntnissen
wird die DCF als Nachbildung eines perfekt fairen Warteschlangensystems
aufgefasst und Abweichungen vom idealen System quantifiziert. Darauf
aufbauend wird ein stochastisches Service-Curve Modell hergeleitet. Dies

ermöglicht beispielsweise die Berechnung von Paketverzögerungen durch die DCF sowie die Anwendung des stochastischen Netzwerkkalküls.

Kapitel 6 präsentiert Ergebnisse einer Messstudie von Methoden zur Bestimmung der verfügbaren Bandbreite in drahtlosen Netzen unter Verwendung der DCF. Es wird gezeigt, dass bekannte iterative Messmethoden, die in der Regel eine first-come first-serve Verarbeitung der Pakete annehmen, den fairen Bandbreitenanteil messen und nicht, wie eigentlich beabsichtigt, die noch verfügbare Bandbreite. Direkte Messmethoden liefern weder den fairen noch den verfügbaren Anteil an der Bandbreite. Wie bereits in vorausgegangen Studien, zeigte sich eine große Abhängigkeiten der geschätzten Bandbreite bezüglich des Messverkehrs und insbesondere der Paketgröße. Im Unterschied zu diesen Studien wird gezeigt, dass bereits kleine Änderungen bestehender Messtools, wie beispielsweise Pathload, ausreichen um auch in drahtlosen Netzwerken zu guten Ergebnissen zu kommen. Des Weiteren wird eine neue Methode zum Abschätzen des fairen Bandbreitenanteils in Funknetzen entwickelt und evaluiert, die eine kontinuierliche Überwachung des Systems erlaubt. Die Arbeit identifiziert die wichtigsten Herausforderungen und analysiert den Trade-Off zwischen einer schnellen Detektion von Änderungen auf der einen, und stabilen Werten auf der anderen Seite. Hierzu wird der Nutzen eines Kalman-Filters, der unter bestimmten Bedingungen optimale Ergebnisse liefert, zum Filtern einzelner Messwerte aufgezeigt und es werden Möglichkeiten diskutiert, das Kalman-Filter optimal zu parametrisieren. Darüber hinaus wird der Einfluss von Messverkehrcharakteristiken auf die Ergebnisse untersucht. Abschließend wird gezeigt wie Anwendungen, wie beispielsweise Video-Streaming, von exakten Messungen verfügbarer Bandbreite profitieren können.

CONTENTS

LIST OF FIGURES

Part I

DISSERTATION

INTRODUCTION

Ever since Hertz discovered radio waves in 1888, and Popow as well as Marconi transmitted radio waves over long distances in 1894, wireless communication has evolved into a complex system with a wide range of different subscribers and applications. In 1970 the first packet-based wireless computer network, ALOHAnet, was developed in ground-breaking work by Abramson [14] at the University of Hawaii to connect seven hosts on four different islands. The basic technology, namely the ALOHA protocol, used in that network has not only built the foundation of wireless but most kinds of modern packet-switched communication networks. Today, it is still used, e.g., in the cellular radio system GSM for its random access channel [139].

Almost three decades later, after several years of development and various generations of proprietary wireless local area network devices, like WaveLAN by Bell Labs/Lucent, the Institute of Electrical and Electronics Engineers (IEEE) ratified its first version of the IEEE 802.11 wireless LAN standard in 1997. Subsequently, in 1999, the WiFi Alliance, a non-profit organization initiated by several major vendors to verify IEEE 802.11 devices in order to ensure their interoperability, was founded. Hence, we use WiFi as a synonym for IEEE 802.11 devices in the latter. With the rise of laptops and other mobile devices in combination with significant reductions in costs, IEEE 802.11 has become the leading standard for regularly used wireless local area networks. Consequently, more than 350 million WiFi devices were shipped in 2010 [68] and today almost every new laptop, tablet pc, and mobile phone comprises a WiFi interface.

To date, the vast majority of WLAN systems have been based on IEEE 802.11 successors such as IEEE 802.11b, IEEE 802.11g, and IEEE 802.11a. These systems basically provide physical-layer based throughput enhancements over the original standard. In 2010, the upcoming IEEE 802.11n

standard that increases the data rate up to 600 Mbps was approved. Among physical-layer enhancements such as the multiple-input multiple-output (MIMO) concept using several antennas and doubling the channel spacing to 40 MHz, the improvements for reaching high data rates are achieved by new MAC layer features such as Block Acknowledgements, Transmission Opportunities (TXOP), and Direct Link Setup (DLS). Besides the well-deployed standard enhancements, there exist other versions for Quality of Service (QoS) support (IEEE 802.11e), transmit power control and dynamic frequency selection (IEEE 802.11h), security and encryption support (IEEE 802.11i), and for vehicular mesh networks (IEEE 802.11p). These standards, partly still under development, are meant to cover even more wireless applications.

The great success of IEEE 802.11 as well as its widespread deployment, however, yields challenges caused, e.g., by a high utilization of the shared wireless medium due to the high number of participants. As a result, the overall goodput of the whole system may decrease due to interferences of the different wireless stations even if the physical data rate of every single participant increases. We conclude that technical possibilities and hardware production costs were crucial at the beginning of wireless networks. Today, however, interoperability, the number of participants, and protocol efficiency are the key factors that limit performance. These challenges will rise in the near future, which motivates the need of networks engineers for a simple, intuitive, and yet accurate framework to analyze, measure, and predict the performance regarding throughput, delay, jitter, and backlog in order to optimize existing and future wireless networks.

1.1   KEY ISSUES IN WIRELESS LAN PERFORMANCE ANALYSIS

In the following we consider a network model as depicted in Figure 1 comprising several wireless nodes within the same sensing and interference range. Thus, every station can overhear the transmissions of all other nodes.

Figure 1: A wireless network model comprising several wireless stations connected to a single access point that acts as a service set coordinator. All stations of interest are within the same interference and sensing range so that they can overhear their transmissions. Thus, they span a single collision domain. This setup is called a *basic service set*. In addition, we define a *tagged station* used for various performance analysis.

In addition, all stations are associated with one access point with whom they communicate with exclusively. The access point can be connected to a wired backbone network. According to the IEEE 802.11 standard, this setup is called a basic service set (BSS) or cell. Throughout this thesis we also use the term *wireless network* as a synonym. In case a station wants to transmit data to a receiver, the nodes cannot establish a direct connection, but the access point has to forward the packets even if the two stations are part of the same BSS. Obviously, this creates a very simple wireless multi-hop network. Without loss of generality, we assume fixed stations even though the node can be mobile in general. Unless otherwise noted, we use the term *wireless link* meaning a wireless broadcast channel with medium access control and radio link control protocols respectively. The network is traversed by packets that can be characterized by different parameters such as their source address, their destination address, their traffic or priority class, and all possible combinations of that. All packets that belong to the same characterization are said to be within the same traffic *flow*.

Compared to modern wired packet-switched systems, wireless networks face some additional challenges which makes their analysis difficult. Re-

garding the physical layer, wireless networks use an unreliable and often unpredictable medium that significantly introduces interference, fading, and path loss which may lead to bit errors and therefore packet loss. Furthermore, wireless networks cannot detect these errors directly, because typically they cannot send and receive data at the same time[1] and check for successful transmissions. This is opposed to wired networks where packet loss occurs mainly due to buffer overflows. The rate of bit errors, which in addition can be detected by the sender, is low enough to be ignored in optimizing the protocols and analyzing their performance. Also related to these issues are the physical transmission paradigms used in wireless networks. Today, most of the transceivers are based on a single-input single-output (SISO) technique. The upcoming IEEE 802.11n standard, however, already uses multiple-input multiple-output streams (MIMO). This increases the physical layer data rate further and is therefore predicted to become the leading technology for wireless networks in the near future. However, making use of several radio streams that affect the performance of a wireless system independently complicates the analysis of these systems. Likewise, wireless links cannot be assumed to be time-independent. Potentially hidden wireless stations, which may even include systems that implement different radio standards using the same frequency band, can join and leave the network at any time, causing unpredictable interferences. Moreover, the environment may change due to movement of obstacles or nodes. These effects can cause rapid fluctuations of the signal-to-noise ratio and may lead to high bit error rates. Different modulation and coding schemes combined with rate adaptation may be used for compensation. As a consequence, the capacity and the availability of the channel may vary drastically. These issues should be considered when designing and analyzing wireless networks. For performance analysis related to the physical layer of wireless channels we refer to, e.g., [57, 113, 142] and references therein.

---

1 In 2010, Choi et al. presented a mechanism that allows for full-duplex transmissions on a single channel at least for low-power SISO systems. For a detailed discussion including the limitations as well we refer to [44] and references therein.

Regarding the data link layer, the medium access procedures of multi-user wireless systems tend to be more complex compared to their wired counterparts due to the aforementioned physical-layer issues and the wireless broadcast medium. In order to deal with the frequent packet loss of wireless links many standards include an automatic repeat request (ARQ) technique such as stop-and-wait ARQ, using explicit acknowledgements and link layer retransmissions to ensure packet delivery. However, those retransmissions consume channel capacity and may cause increasing and varying one-way delays. Hence, they have to be taken into account when analyzing the performance of wireless systems. Furthermore, in case of multi-access systems, stations that share the same medium need to contend for channel access. The contention phase is controlled by the MAC protocol, often in a fully distributed manner. This is already true for wired MAC protocols, like IEEE 802.3 Ethernet, that have originally been designed to work in broadcast environments as well. However, today these systems are typically connected by store-and-forward switches and routers in a point-to-point manner. This renders the contention phase for a channel access unnecessary but calls for a specific packet handling at the intermediate network systems. Implementing similar point-to-point connections in wireless systems, e.g., by frequency multiplexing, is not suitable since it would limit the number of stations dramatically. At the end of the day, the analysis of packet scheduling implemented in wired network devices comes to an analysis of medium access procedures in wireless systems. In both cases, the sharing of resources by packets of different nodes or flows is also called *statistical multiplexing*. For wireless systems the behavior may differ largely compared to the multiplexing of flows in a point-to-point infrastructure. For instance, on top of a different scheduling behavior which is explored in Chapter 2 and Chapter 4, the outgoing link in wireless multi-hop networks may also affect incoming traffic since it also occupies the shared medium. Moreover, these characteristics of wireless links are of vital importance

for applications such as available bandwidth estimation, as discussed in Chapter 6.

In addition to the wireless specific challenges, we have to deal with the key issues that generally arise in packet-switched networks. For instance, nodes require the availability of *buffers* to store packets temporarily in case that the amount of packets to be served exceeds the channel capacity only for a very short period of time. To this end, packets can experience additional queuing delays, which also affect the end-to-end performance of flows. Other challenges stem from the *traffic characteristics* that may exhibit high variabilities or correlations and need to be modeled accurately.

All these challenges mentioned above make the performance evaluation of wireless packet-switched networks difficult.

1.2   THEORIES FOR WIRELESS NETWORKS PERFORMANCE ANALYSIS

In this section we provide a brief overview of analytical and empirical methods to analyze performance measures such as throughput, delay and backlog in wireless networks.

The problem of network performance analysis in general has a long history. For wired communication networks, it dates back to 1909 and 1917 when Erlang published his seminal work on *classical queuing theory* [55, 56] that allows the calculation of various performance metrics such as average waiting time, expected number of entities, and blocking probabilities of limited queues and buffers. Among many other applications this theory has been used for the analysis and the dimensioning of telephone systems. In 1962, Kleinrock published his ground-breaking work that proves the superiority of packet-switched networks over circuit-switched network in terms of resource utilization and efficiency using classical queuing theory [91, 95]. Subsequently, queuing theory delivered important insights such as calculations for delay, throughput, and backlog of various systems into the field of wired computer networks. In its very basic form, the output

link of network components such as switches and routers can be modeled by an $M/M/1$ queue that generally represents a single server with an unlimited buffer capacity, Markovian arrival process, and exponentially distributed service times. Traffic flows that traverse multiple nodes can be modeled as a sequence of concatenated $M/M/1$ queuing systems that are known to behave as independent $M/M/1$ queues [32, 46, 69]. Regarding the state probabilities, the steady-state solution of such systems describes the system's behavior in an unambiguous way and can be calculated as a product of the steady-state solutions of each of the queues. These systems are also called *product-form networks* [32].

Likewise, the underlying theory of *Markov chains* that basically constitutes the fundamental concept of queuing systems, has been used to model medium access procedures in order to derive performance results including throughput and access delays, for wireless networks. In general, Markov chains are Markov processes, i.e. first-order dependent stochastic processes, with a discrete state space. Thus, the probability of reaching a given state only depends on its previous state, but is conditionally independent on any other state in the past. Furthermore, if the Markov chain is time-homogeneous, state transitions do not depend on the observation time and any successive state does not depend on the time the process has already spent in the current state. The process is memoryless, which implies geometrically distributed state waiting times for time-discrete systems and exponentially distributed waiting times for time-continuous systems respectively. For further information and a detailed introduction to Markov chains in general we refer to [32]. An application of Markov chains to model the IEEE 802.11 DCF can be found, e.g., in [31].

Regarding its characteristics, however, the fundamental assumption of memoryless packet inter-arrival and service times that yields a Poisson process does not model Internet data traffic accurately since that kind of traffic exhibits correlations and long-range dependencies that yield self-similar properties, at least in wired networks. This is shown in the seminal

work by Leland et al. [103] published in 1994. These findings have been
supported subsequently for various types of network traffic including
wide-area-network traffic [110, 121], world-wide-web traffic [50], and video
traffic [26] and obviously limit the application of queuing networks for
analyzing computer communication networks. The objective to overcome
this problem motivates research of sophisticated traffic models like frac-
tional Brownian motion that is meant to describe the behavior of Internet
traffic accurately [103, 117, 118]. Furthermore, new approaches such as effec-
tive bandwidths and deterministic network calculus have been developed
aiming at a general methodology for network performance analysis.

The theory of *effective bandwidths*, actually developed for ATM networks
and widely used for an efficient analysis of statistical multiplexing, describes
the minimum required bandwidth that must be provided to a flow so that it
can meet certain service guarantees. The effective bandwidth, described by a
space and a time parameter that contain information about the distribution
of the flow and the length of a time period respectively, is located between
the average bandwidth and the peak rate of a flow. Effective bandwidths
are known for a wide range of arrival processes. Moreover, the effective
bandwidth of aggregated flows can be determined by simply summarizing
the effective bandwidths of the individual flows. However, the theory of
effective bandwidth covers just a few scheduling models and its application
is limited to asymptotic approximations. For a detailed introduction to
effective bandwidths we also refer to [43, 88] and references therein. For
extensions covering the analysis of wireless networks see, e.g., [104].

The *network calculus*, developed in the pioneering work by Cruz [51],
evolved from the calculus for network delays and the concept of determin-
istic service curves for Generalized Processor Sharing (GPS) schedulers
proposed in [120]. Using min-plus algebra, a queuing system, as depicted
in Figure 2, is described by its traffic arrivals $A(t)$ cumulated over time at
the input, the systems service, and the cumulative traffic departures $D(t)$
at the output. Typically, the input is described by an envelope function

Figure 2: A queuing system that links the arrivals $A(t)$ of a system to its departures $D(t)$ by a service curve $S(t)$.

covering the worst case of traffic arrivals, whereas the service of a system is characterized by a service curve $S(t)$ that represents a lower bound of an amount of service that is offered to a flow. Similar to systems theory, a service curve represents the impulse response of a linear system, however, under min-plus algebra and all metrics are linked by

$$D(t) = \inf_{\tau \in [0,t]} \{A(\tau) + S(t - \tau)\} =: A \otimes S(t)$$

where $\otimes$ is known as the min-plus convolution. To this end, the departures can be calculated from the arrivals and a service curve representation, which is known for various network elements such as switches, routers, and links and their corresponding schedulers. Furthermore, the convolution function is associative and a whole system of concatenated network elements can be represented by the convolution of all single service curves along a network path. The end-to-end service curve, that basically reduces the many-node case to a single-node case that describes the end-to-end available service, follows readily. Due to the worst-case representation the deterministic network calculus can be used for strict service guarantees. It does, however, not incorporate any statistical multiplexing gain and frequently yields overly pessimistic performance bounds, even if the worst case is known to happen very rarely. Additionally, most of the real-time applications such as voice, video, and online games, tolerate moderate quality of service violations [39] and still achieve acceptable results. Hence, a worst-case performance bound could be relaxed in order to save resources without a significant reduction of quality experienced by the user. Moreover, there are various systems, like wireless radio channels, that are unable to provide any strict service

guarantees and, thus, cannot be modeled by the deterministic network calculus.

To overcome the limitations of deterministic network calculus and to incorporate the statistical multiplexing gain, recently, significant progress has been made towards the formulation of stochastic service curves, see, e.g., [59, 80] and references therein. Service guarantees that can be calculated using stochastic network calculus are generally described by percentiles

$$P\left[metric > x\right] \leq \epsilon$$

that state the probability a given metric, such as delay and backlog, exceeds a certain value x is less or equal a certain violation probability ϵ, see, e.g., [33]. Usually, the violation probability is quite small, e.g., in the order of 10^{-6} to 10^{-9}, to exclude only some rare events. These small values can improve or, like for wireless networks, make it possible to provide network performance bounds. Stochastic network calculus models are used, e.g., in [57, 113], to derive service curve representations of wireless links with a focus on channel outages that are due to fading and interference. Modeling random medium access and quantifying deviations of the DCF from the fair share on different time scales, however, remains an open issue that is addressed in this work.

Closely related to the min-plus algebra typically used by network calculus is the max-plus algebra, where the service curve is inverted from a function of time to a function of packets, see, e.g., [21]. Regarding the deterministic network calculus, this relation establishes a close link between the Guaranteed Rate Clock model [61] and so-called latency-rate service curves. Furthermore, in max-plus algebra we can operate on packet time stamps and cover packetized traffic directly which simplifies the analysis of real systems. In Chapter 5 we use max-plus algebra to phrase our model.

1.3 THESIS STATEMENT AND CONTRIBUTION

This thesis attempts to shed light on the interaction of various wireless nodes. It demonstrates how a better understanding of the interaction, i.e., the fairness throughout wireless stations, can lead to better performance evaluations and measurement methodologies. Hence, we phrase our overall idea by the following statement.

> *A solid understanding of the interaction of various wireless stations, that is brought down by new analytical insights, leads to better performance evaluations and measurement methodologies.*

This thesis contributes several theoretical, empirical, and practical results. We develop a stochastic service curve model for IEEE 802.11 wireless networks, show its accuracy and applicability, and present application that can benefit from such a model. We start by analyzing the fairness of the DCF at a single radio channel without hidden or exposed terminals. We derive closed-form solutions for the conditional distribution $P[K{=}k|l]$ that $M-1$ contending stations transmit k packets given a tagged station M transmits l packets. This characterization of fairness turns out to be comprehensive and versatile, e.g., the well-known fairness index by Jain [75] follows readily. We substantiate our analytical findings using an extensive set of measurements that were partly conducted in a shielded room to avoid uncontrollable external influences as well as OMNet++ simulations. Moreover, we view the DCF as emulating the Generalized Processor Sharing policy. We formulate a recursive model for packet departure times coined the DCF clock, which is subject to well-defined random error terms. Based on the distribution of packet inter-transmissions we derive a stochastic service curve model for the DCF and use this model to calculate probabilistic service bounds. To verify the results we apply external measurement techniques, originally provided for wired-router performance evaluation, which make use of network calculus models and parameters such as packet rate and latency. We show

how the proposed methods can be applied to wireless networks and present extensive measurement and simulation results quantifying the performance of wireless networks. Using our framework, we show how user applications, such as rate adaptive video streaming, can estimate their fair bandwidth share under the DCF from passive measurements of their data arrivals and departures. Closely related are active probing techniques that seek to identify the unused capacity, also referred to as available bandwidth, along a network path from specific probing packets.

1.4   THESIS STRUCTURE

The remaining part of this thesis is structured as follows.

In Chapter 2 we provide a brief introduction to the IEEE 802.11 standard and related MAC mechanisms. We then summarize existing work on performance analysis and related measurement methods. Finally, we review related work on fairness issues with a focus on wireless networks.

Chapter 3 refers to the challenges that are addressed in this paper and lists the main research questions. It justifies that these questions are previously unanswered. Moreover, it motivates why it is worthwhile answering these questions.

In Chapter 4 we investigate the long- and short-term fairness of the IEEE 802.11 medium access procedure. Here we present some of the main results of this thesis, i.e. a fairness model that describes the fairness of the DCF achieved over various time scales and various number of hosts. We substantiate our analytical findings by presenting an extensive set of simulations and measurements.

In Chapter 5 we formulate a stochastic service curve model that is suitable for analyzing the performance of wireless networks using the DCF. We use this model to investigate performance bounds of the wireless network. Moreover, we compare the results to simulations.

In Chapter 6 we present applications, such as available bandwidth estimation and adaptive video streaming, that can benefit from the previous results. We investigate existing methods for active available bandwidth estimation and present theoretical foundations to bandwidth estimation in wireless systems. Moreover, we present our own measurement methodology based on our findings in Chapter 4. To this end, we apply Kalman filtering to accurately track varying available bandwidth over time. Moreover, we show how the Kalman filter can be tuned optimally.

In Chapter 7 we present conclusions and future work.

2

BACKGROUND AND RELATED WORK

In this chapter we introduce the IEEE 802.11 standard and the basic functionalities of its medium access control mechanisms. We discuss related work on its performance and present basic results derived from that. Furthermore, we survey related work regarding fairness in infrastructured wireless networks based on IEEE 802.11 and provide a brief outlook on applications, such as active bandwidth measurement techniques, which are affected by these fairness issues.

2.1 THE IEEE 802.11 STANDARD

Due to the good performance and great success of the IEEE 802.3 Ethernet standard [12] that defines the physical (PHY) and medium access control (MAC) specifications for wired local area networks, IEEE 802.11 [11] is based on the same distributed medium access mechanism named carrier sense multiple access (CSMA). With CSMA in general, a station that intends to transmit a packet listens to the shared medium for a given amount of time. If the medium is idle during this period, the station transmits its packet. Since several stations share a single medium to transmit packets, collisions may occur if more than one station transmit data at the same time. In order to deal with this problem, the original Ethernet standard uses CSMA with collision detection (CSMA/CD) where a station receives its own transmission and detects collisions by comparing the transmitted to the received signals. In case of a collision, all colliding stations back off for a uniformly distributed random period in a given interval $[0, u]$, that is also called collision window. This random backoff period reduces the probability of a second collision. However, if another collision occurs anyway, the colliding stations double their collision window at most up to a
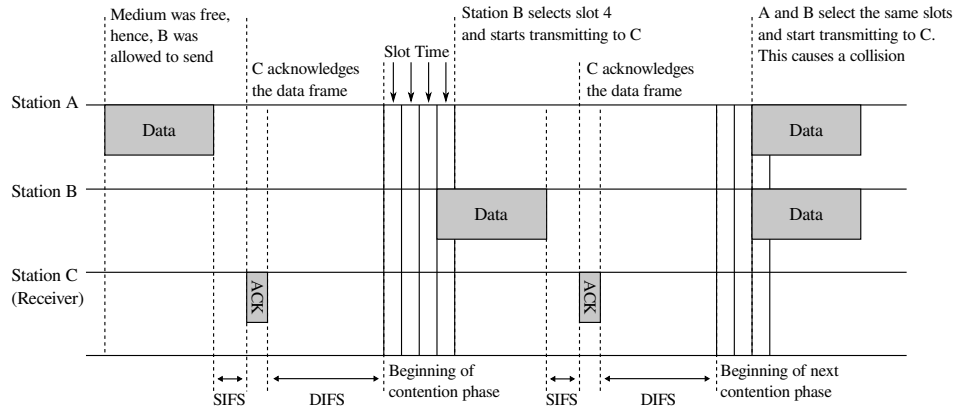
Figure 3: The IEEE 802.11 distributed coordination function using a carrier sense multiple access procedure with collision avoidance in case of three competing stations. Additionally, slot times and the various interframe spaces, i.e. SIFS and DIFS are highlighted [11].

maximum value u_{max} and return to u_{min} in case of a successful transmission. Furthermore, if the number of retransmissions reaches an upper limit the corresponding packet is discarded. Since in the range from u_{min} to u_{max}, the actual collision window size u meets the exponential distribution $2^x - 1$, this procedure is called binary exponential backoff.

Using CSMA/CD in wireless networks is not suitable since interferences become relevant only at the receiver. At the same time, it is not possible for a sender to detect this interferences or collisions by listening to the channel while transmitting. WiFi equipment typically supports half-duplex transmissions, i.e. communication in both directions, but only in one direction at a time, mainly for historical and cost reasons. Moreover, even if the equipment would support full duplex transmission, i.e. communication in both directions simultaneously, it may miss collisions anyway if they occur outside the senders sensing range. Therefore, the distributed coordination function (DCF) of IEEE 802.11, that is the fundamental access method used to support a data transfer on a best effort basis, uses carrier sense multiple access with collision avoidance (CSMA/CA) in combination with explicit acknowledgements (ACKs) for all packets to detect collisions as well as transmission failures in general. As shown in Figure 3, we consider a tagged station A that transmits a packet and receives the

corresponding acknowledgement after the so-called short inter-frame space (SIFS) which is a hardware constant given by the corresponding PHY implementation. Before sending the next packet the station generates a uniformly distributed backoff variate b in an interval $[0, w-1]$. This interval is also called *contention window* of size w. Related to IEEE 802.3 a station that intends to transmit a packet performs a carrier sensing to detect the presence of other WiFi stations. Whenever the medium is continuously idle for the duration of a distributed inter-frame space (DIFS) that is calculated by *Slot-Time* $+ 2\,SIFS$, the station starts, respectively continues, its countdown procedure. The backoff value is decremented by one after each PHY-specific slot time. This countdown is paused if the medium is sensed to be busy again and it is resumed after each successful DIFS waiting. As soon as the backoff value reaches zero the station sends its packet. If several stations perform the countdown procedure simultaneously the one with the smallest backoff value starts its transmission first. If two or more stations finish the countdown procedure at the same time their transmissions cause a collision. In this case the respective stations perform an exponential backoff as in the wired case, i.e. they double their current contention window w at most up to w_{max} and return to w_{min} in case of a successful transmission. Doubling the collision window, however, has historical reasons since this procedure has been easy to implement using a binary shift operation. Yet, performing a more complex adaptation procedure, like multiplicative increase linear decrease (MILD) and additive increase multiple decrease (AIMD), may result in a better network performance. In any way, the collision avoidance algorithm adapts to the number of collisions which is related to the number of contenting stations[1]. Additionally, to improve the performance in case of moderate channel load, the DCF performs a post-backoff procedure immediately after a packet is sent. Thus, after a long time not performing any

---

1 Actually, the collision avoidance mechanism adapts to transmission failures in general that are assumed to be collisions. However, in a real-world scenario this might not be the case. For a more detailed analysis of transmission errors we refer to, e.g., [60, 90] and references therein. Solutions proposed to distinguish between transmission errors and collisions can be found, e.g., in [90]

transmission attempt, the backoff counter becomes zero and a station that has already performed such a post backoff can send a new frame directly after DIFS period. For a more detailed introduction to IEEE 802.11 MAC mechanisms including the DCF we refer to [11, 122, 143].

To cope with the hidden station problem that is explored, e.g., in [140], the IEEE 802.11 MAC may use a three-way handshake called request-to-send/clear-to-send (RTS/CTS). It reserves the channel using small control packets before sending its actual data packets in order to minimize the costs in case of collisions to the expense of an additional overhead. Since the RTS/CTS handshake is often disabled by default in today's network configurations, we neglect it in our following analysis. However, even if RTS/CTS is enabled the basic DCF mechanism is still in charge to grant channel access. Thus, our performance analysis can be easily transformed to a case incorporating RTS/CTS.

In addition to the standard DCF, IEEE 802.11 provides an optional centralized channel access mechanism called point coordination function (PCF). It resides in a point coordinator, which typically is the access point, and coordinates the channel access of each station by additional control frames. The point coordinator waits for a PCF inter-frame space (PIFS) which is smaller than DIFS but larger than SIFS and provides some kind of priority to grasp the channel and send its data. Obviously, this allows wireless networks to support some QoS guarantees. However, the PCF seems to be implemented only in some wireless devices maybe because it is not part of the WiFi interoperability standard. Thus, we neglect PCF driven wireless networks in this thesis.

Using the given procedures and standard values of the IEEE 802.11 MAC and PHY specifications [11], it is possible to derive several basic performance results: the effective throughput, also known as goodput, of a single station as a function of packet sizes, the collision probability as a function of the number of stations and mean contention window size, and

Characteristic	DSSS 802.11	HR-DSSS 802.11b	OFDM 802.11a	ERP-OFDM 802.11g	DSSS-OFDM 802.11g
$datarate_{max}$	2 Mbps	11 Mbps	54 Mbps	54 Mbps	54 Mbps
slot-time	20 μs	20 μs	9 μs	9/20 μs	20 μs
SIFS	10 μs	10 μs	16 μs	10 μs	10 μs
DIFS	50 μs	50 μs	34 μs	28 μs	50 μs
ACK	14 Byte	14 Byte	14 Byte	14 Byte	14 Byte
CW_{min}	31	31	15	15/31	31
CW_{max}	1023	1023	1023	1023	1023
$Preamble_{PLCP}$	144 μs	144/72 μs	16 μs	16 μs	144/72 μs
$Header_{PLCP}$	48 μs	48/24 μs	4 μs	4 μs	48/24 μs
$Preamble_{OFDM}$					8 μs
$Header_{OFDM}$					4 μs
signal-extension				6 μs	6 μs
$Symbol\,Interval$			4 μs	4 μs	

Table 1: Details regarding transmission times in 802.11a/b/g. [11]

the contention window size as a function of the number of stations and collision probability.

IEEE 802.11 specifies various timing values for the different transmission modes and coding schemes related to physical-layer features and capabilities. A subset of theses values that reflect the most common and therefore most relevant schemes are shown in Table 1. Starting from its basics, IEEE 802.11 transmits in the 2.4 GHz industrial, scientific, and medical (ISM) band with 20 MHz bandwidth using a direct-sequence spread spectrum (DSSS) modulation technique [143] that supports data rates of 1 and 2 Mbps for the payload and 1 Mbps for the physical-layer convergence procedure (PLCP) preamble and header respectively. Subsequently, IEEE 802.11b transmits payload with data rates up to 11 Mbps using a so-called high-rate (HR) DSSS. The PLCP preamble and header are, however, transmitted at 1 Mbps as in the basic IEEE 802.11 case. Furthermore, IEEE 802.11b optionally supports short preamble and header to reduce the overhead and, hence, increases the throughput. The IEEE 802.11a standard transmits in the 5 GHz band typically with 20 MHz bandwidth using an orthogonal frequency-division multiplexing (OFDM) modulation technique [143] that supports data rates of 6 to 54 Mbps for the payload. Since 802.11a does not

need to be backward compatible to previous standards, due to its operation in a different frequency band, it provides new PLCP preamble and header formats that are transmitted at higher speed and, thus, further reduce the overhead. To support data rates up to 54 Mbps also in the 2.4 GHz band, IEEE 802.11g introduces the extended-rate physical (ERP) layer that combines the transmission procedures of the previous IEEE 802.11 standards in 2.4 GHz with the OFDM coding scheme of IEEE 802.11a. In case only ERP stations are associated with an access point, ERP-OFDM achieves the very same performance results in terms of achievable throughput as in the 5 GHz band. If there are IEEE 802.11b devices connected as well, the timing parameters are changed, e.g., using DSSS-OFDM, to assure backward compatibility and interoperability of all these devices. This, obviously, reduces the performance.

Next, we evaluate the impact of IEEE 802.11 timing parameters and DCF protocol overhead on the achievable throughput. Given a frame size of L bytes, the long-term-achievable throughput, that is also referred to as goodput, can be determined by $G = L/g$. Here, the gap g denotes the mean time between the beginning of two subsequent frames transmitted at maximum rate. This gap consists of a number of additive parameters such as inter frame spacings (t_{difs}, t_{sifs}), the expected backoff time ($t_{backoff} = E[CW_{min}]$ $slot\text{-}time$), transmission times ($txtime$, t_{ack}) for data frames and acknowledgements including preamble and header, and propagation delay that, in addition, is neglected here, since it is specified by the standard to be smaller than 1 μs. We compute the transmission time $txtime(L)$ using the formulas provided by the standard and derive the maximum achievable throughput subsequently for specific transmission modes in an optimal single node case as follows

$$G = \frac{L}{t_{difs} + t_{backoff} + txtime(L) + t_{sifs} + t_{ack}}. \tag{2.1}$$
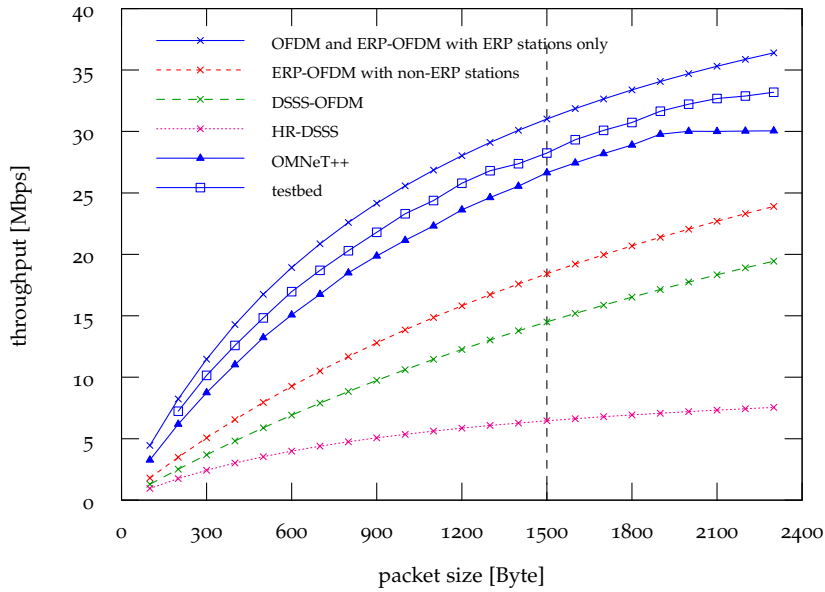
Figure 4: Packet size dependent throughput in IEEE 802.11 for different transmissions schemes derived from an analytical model. We compare the analytical findings to results from simulations and measurements respectively and find the analytical model matches the shape of the empirical data well.

We depict the analytical results for different packet sizes and different transmission schemes in Figure 4. Regarding the measurement and simulation results that are gathered from the testbed as described in Chapter 4, we calculate the mean of 25 experiment runs per packet size lasting for 60 seconds each. We present the outcome without showing the 0.95 confidence intervals since they are negligible small. By comparing the analytical and empirical results we find that they match quite well and the analytical model provides an upper bound for the throughput in wireless LANs. Moreover, we conclude that the size of packets and the traffic characteristics have a large impact on the performance of wireless networks. This has to be considered for various applications such as available bandwidth estimation and its probing packets.

2.2   RELATED WORK ON PERFORMANCE EVALUATION IN IEEE 802.11

In the sequel, we elaborate and discuss related work on the collision probability, the distribution of backoff values as a function of contending stations, and the corresponding throughput. As throughout this thesis, all models considered here, focus on the MAC part of a wireless system and we assume ideal channel conditions.

For pure random access wireless systems like ALOHA, where hosts send their data independently, i.e. without taking other hosts into consideration, Abramson provided a first and general performance analysis regarding the averaged throughput as a function of offered channel load that is affected by frame collisions [14, 15]. Assuming a large number of sending stations $M \to \infty$, a fixed probability τ that a station transmits at an arbitrary point in time, and constant sized frames with a transmission time T_t normalized to 1, it can be shown that the number of frame arrivals at the channel is Poisson distributed and the probability of simultaneous packet arrivals in a given time period can be calculated. A frame is successfully transmitted if and only if there is just one arrival in an interval of length $2T_t = 2$ so that no collision occurs. Thus, the average throughput becomes $S = \lambda e^{-2\lambda}$, where λ is the average frame generation rate. The maximum throughput follows readily as $S_{max} = 1/(2e) \approx 0.18$. In [126] Roberts et al. present an intuitive way to make pure ALOHA more efficient by using a slotted time algorithm that allows nodes to transmit their packets only at slot boundaries. Furthermore, the reception of a packet is completed within one slot. Thus, the period where a collision might occur reduces from $2T_t$ to T_t and the throughput becomes $S = \lambda e^{-\lambda}$. Likewise, the maximum throughput is doubled and follows by $S_{max} = 1/e \approx 0.37$. For a summery of the ALOHA protocol, its application, and more detailed information we refer to [95] and references therein. Further results, such as for retransmissions, stability, and access delays can also be found, e.g., in [29, 87, 92].

In 2000, a very similar approach using a p-persistent protocol model for IEEE 802.11 that samples its random backoff interval from a geometric distribution with parameter p was proposed by the seminal work of Cali et al. [41]. Thus, the probability of transmission attempts per idle slot remains constant over time and the sequence of collisions and successful transmissions is independent and identically distributed. Knowing the average channel access time and tuning the geometric backoff distribution accordingly, it has been shown by the authors that this approach delivers accurate average throughput and delay results.

At the same time, another pioneering performance analysis regarding the saturation throughput of IEEE 802.11 and the DCF was provided by Bianchi [31]. To tackle the challenging task of describing the behavior of the binary exponential backoff algorithm that couples the service processes of various contending stations throughout their shared collisions, the fundamental step here is the decoupling approximation. This states that every node has a given and constant collision probability p which is independent of the history of the node itself as well as of all other stations respectively. Despite other differences, this is very similar to the ALOHA analysis and the p-persistent model presented above where the transmission attempt probability τ has the very same properties yielding a steady-state collision probability $p(\tau) = 1 - (1-\tau)^{M-1}$ with M contending station. For Bianchi's DCF analysis, $\tau(p) = \sum_{i=0}^{m} b_{i,0}$ follows from a two-dimensional Markov chain, where m is the maximum retransmission counter and $b_{i,j}$ are the state probabilities of all backoff states that are closely related to the actual backoff values a station can sample. The author proves that the resulting non-linear system of equations that comprises $\tau(p)$ and $p(\tau)$ yields a unique solution and can be solved numerically. By making use of the calculated results, other network performance metrics such as average long-term throughput can be calculated. Throughout a comparison to simulations, the results have been shown to be quite accurate also for setups comprising a small number of contending stations. Figure 5 depicts the results for the average collision

(a) average collision probability   (b) average backoff value
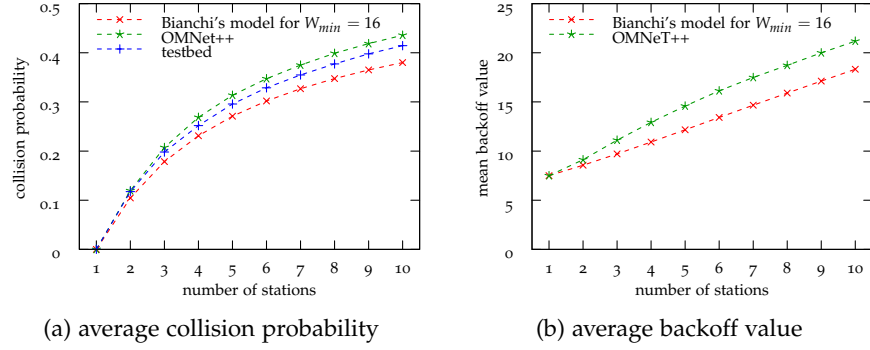
Figure 5: Applying the decoupling approximation introduced by Bianchi [31] and Cali [41] respectively, we calculate the mean collision probabilities and the related average backoff values over the number of stations. The theoretical results match testbed and simulation data quite well.

probability and the mean backoff samples. We compare the model outcome to simulation and testbed results and confirm a good matching also for a limited number of contending stations.

Subsequently, Bianchi's approach is approximated using a one-dimensional Markov chain in [96] and further simplified using backoff values sampled from an exponential distribution with a mean related to the collision probability in [95], respectively. By using Bianchi's approach, channel access delays are derived, e.g., in [42, 96, 129].

Owing to the related work discussed above, we conclude that modeling the backoff procedure of IEEE 802.11's DCF using exponentially distributed backoff samples yields accurate performance results for average throughput and delay calculations.

Complementary to analytical performance evaluations, many studies attempt to analyze the performance of networks by measurements. For a general overview we refer to, e.g., [71, 74, 89, 128] for wired and to [16, 22, 146] for wireless systems, and references therein. In the following, we focus on a subset of performance metrics, namely the available bandwidth, especially in wireless networks. The term available bandwidth denotes the portion of the capacity at a link or a network path that remains unused by present traffic. The idea to estimate the bandwidth of a network path

from end-host measurements dates back to TCP congestion control [70] and packet pair probing [89]. Since then, the field of available bandwidth estimation has evolved significantly and to date a number of estimation methods exists, e.g., [67, 71, 114, 125, 137], which are frequently used, e.g., for network management, error diagnostics [66], overlay routing [73, 149], traffic engineering [94], and video adaptation [13, 116]. The theoretical underpinnings of bandwidth estimation have been explored, e.g., in [107, 108, 109, 112] and empirical evaluations can be found, e.g., in [135, 137]. Common challenges and pitfalls regarding the design and evaluation of bandwidth-estimation tools can be found in [72].

The task of bandwidth estimation in wireless networks, such as IEEE 802.11 Wireless LANs, however, has been understood to a much lesser extent. Bandwidth-estimation methods, which perform well in case of wired links, have been reported to yield highly unreliable available bandwidth estimates for wireless links [34, 98, 105], hinting at a number of specific challenges and open issues in wireless bandwidth estimation. Here, the performance and the quality of service of a link depend largely on the characteristics of the shared physical medium and the multi-access coordination function. These aspects strongly influence quantities like delay, loss, and throughput and may result in a high variability of available or actually accessible resources. This makes measurement-based bandwidth estimation in wireless networks a complex and difficult task.

A few approaches to bandwidth estimation specifically address the characteristics of wireless networks. Passive methods, which measure existing traffic, can take advantage of the wireless broadcast medium and record idle periods to estimate the resources that would be available in the proximity of a node [102, 130]. The approach is, however, unreliable in case of hidden stations. Active probing, on the other hand, takes measurements of specific probing traffic at the ingress and the egress of the network to infer the available bandwidth of a network path. Contrary to wired networks a strong impact of packet sizes on bandwidth estimates has been observed for

wireless links [17, 82, 83, 98, 105]. Here, the fluid model, which is employed by many estimation methods, is clearly violated. Moreover, as explored, e.g., by [98], the assumption of first-come first-serve (FCFS) scheduling, which is the basis of most active probing methods, does not hold in IEEE 802.11 WLANs. To this end, the scheduling, i.e. fair scheduling in case of the DCF, needs to be taken into account, see, e.g., [38, 124]. Hence, we provide an overview of fairness aspects regarding IEEE 802.11 wireless systems in the following.

2.3 RELATED WORK ON FAIRNESS IN IEEE 802.11

As mentioned before, an important aspect of communication protocols in general is their attainable quality of service. Thus, achieving good throughput while taking a – possibly weighted – fair medium access of an unknown number of users into account is a major goal of MAC protocols. For instance, fairness on long time scales ensures a certain average throughput, while the issue of short-term fairness has tremendous impact on individual packet delays [93, 96, 129]. This not only has a severe influence on elastic transport protocols such as TCP [70, 147], but also on many measurement applications like available bandwidth estimation as explored, e.g., in [38, 124]. Furthermore, it is desirable to have some kind of service guarantees, such as maximum delay or minimum jitter especially for time-critical applications such as voice, video, or games [19, 39, 144]. These challenges gave rise to significant research and for a survey we refer to, e.g., [132]. Using a distributed scheduling approach, such as the DCF, addressing the problems efficiently is, however, still challenging.

In the following discussion we focus on the MAC part of wireless networks and its impact on fairness. Hence, we assume all stations being part of a wireless basic service set, i.e. all stations using the normal DCF mode of operation in infrastructure mode communicating with an access point that serves as a service set coordinator. Furthermore, all stations are within the

same interference and sensing range, spanning a singular collision domain. Since nowadays, wireless networks usually act as access networks to a wired backbone network, like the Internet, this assumption seems reasonable. For works on physical-layer-related issues regarding fairness, including pathloss and asymmetric links, we refer to, e.g., [60] and references therein. Works on infrastructure-less systems, such as wireless ad-hoc or wireless multi-hop networks, can be found, e.g., in [24, 85].

As noted, e.g., in [27, 28], the DCF seeks to achieve *per-packet fairness*, i.e. all stations transmit the same number of packets on average over sufficiently long time scales. This is maintained by the channel contention phase granting access to the channel with equal probabilities for all stations in a distributed manner after a DIFS waiting. In [64] Heusse et al. discovered a performance anomaly of wireless networks using the per-packet fair DCF when supporting various bitrates among different stations. Hosts that occupy the channel for a long time because of using a low physical bitrate penalize other hosts transmitting at higher rates. This results in an overall low throughput approximately at the data rate of the weakest sender. Thus, the notion of *airtime fairness* used, e.g., in [20, 77, 84], that aims to provide a fair channel access not on a per-packet basis but on channel occupation times, becomes relevant in IEEE 802.11. This depends, however, on further parameters such as packet sizes and rate-adaptation and becomes notable especially in wireless networks supporting high but also varying physical data rates. Hence, some of these techniques developed in this field, like TXOP, have also been applied to the IEEE 802.11e and the upcoming IEEE 802.11n standard not only because of fairness issues but to overcome the disproportion of PHY- and MAC-layer overhead and actual transmitted data that occurs mainly due to compatibility reasons. As a result, it is possible to transmit more than one frame per granted channel time and the application of per-packet fairness may become less important in the future. We note, however, that not even the degree of per-packet fairness that is achieved by the standard DCF is well understood and precisely modeled. Moreover, we

believe that deeper insights in these aspects lead to a better understanding as well as better solutions also in the case of airtime fairness.

Fairness is often quantified using Jain's fairness index [75]. Given a number of n samples, k_i it is defined by the first and second moment of the distribution K of k as

$$J = \frac{\mathsf{E}[K]^2}{\mathsf{E}[K^2]} \;=\; \frac{E[K]^2}{Var[K] + E[K]^2} \;=\; \frac{(\sum_{i=1}^n k_i)^2}{n \sum_{i=1}^n (k_i)^2} \tag{2.2}$$

where larger values of $f \in [0,1]$ indicate better fairness. For instance, in a multiuser network system we may use the mean and the variation of a given set of user throughputs k to calculate Jain's fairness index regarding the bandwidth share. If every user achieves the very same throughput the index is equal to 1, whereas it is $1/M$ if only one user out of M occupies the whole channel. Thus, the metric provides intuitive fairness values and can be used to compare the fairness of different systems. However, since Jain's index only provides a singular value that is derived from the first and second moment, it may lack some more complex aspects and insights of fairness in a system.

Actually beyond the goal of providing a general definition of fairness, Bharghavan et al. presented one of the first papers that point at its importance in wireless networks [30]. The authors investigated bandwidth allocation of the pre-802.11 multiple access with collision avoidance (MACA) protocol proposed by Karn [86] and redefined by Biba [10] to be used with Lucent WaveLAN-cards. In line with [133], they find that using the binary exponential backoff procedure implemented by the MACA protocol, i.e. with no maximum backoff, one station eventually occupies the channel permanently under high load by having a significantly lower backoff counter than the others. To cover this, they proposed a MACA extension called MACA for wireless (MACAW) that, among other things such as explicit acknowledgements and a MILD backoff mechanism to improve the performance, shares the backoff counter values using a packet header field. Thus,

all stations within a basic service set have the same backoff counter and use the same backoff distribution, which ensures a somehow fair medium access on long time scales. Complementary to the work presented above, an empirical study of fairness in WaveLANs is provided by Koksal et al. [93]. The throughput of individual stations is averaged over different window sizes to compute Jain's fairness index for short and long time scales. The authors state that WaveLAN systems are long-term fair but unfair on short time scales which basically supports the findings presented in [30].

Contrary to the MACA and WaveLAN protocol, IEEE 802.11 uses a bounded binary exponential backoff mechanism, that alleviates unfairness to some extent to the expense of collisions and packet losses. Similar results are reported for WiFi systems by a number of subsequent papers, such as [24, 97, 115, 141]. Some of them also compute Jain's fairness index from simulation or measurement data. Furthermore, based on a Markov model the study [106] supports short-term unfairness.

In contrast, the authors of [27, 28] reported short-term fairness for IEEE 802.11 based on measurements and an analytical model that introduces a new indicator of fairness. To this end, the packets that are transmitted by contending stations before a tagged station transmits a single packet are counted. In the sequel we refer to these packets as *inter-transmissions*. For instance, we assume two independent, greedy wireless stations M_i with index $i = 1, 2$ which intend to send their frames at maximum speed and contend for a channel with a limited capacity. Due to the random medium access, on the channel this yields a data stream with a random packet order, like, for example: 1 2 2 1 1 2 1 2 2 2 1 1 2 1 2 1. If we tag station M_1 and count the packets of M_2 that are inter-transmitted between two subsequent frames of M_1 we obtain the following number of inter-transmissions k: 2, 0, 1, 3, 0, 1, 1. The distribution K of the random inter-transmissions k is directly related to the fairness of the system. Assuming that $E[K] = 1$, the smaller the variance of K, the better the fairness. To model the distribution K, the authors consider independent and identically distributed (i.i.d.)

random backoff values $b_i(j)$ of packet j at station i. Thus, the first station sends exactly k packets before the second station completes its countdown if $\sum_{j=1}^{k} b_1(j) \leq b_2(1)$ and $\sum_{j=1}^{k+1} b_1(j) > b_2(1)$. The backoff values are modeled as continuous and uniformly distributed in $[0, w-1]$, where w is the minimal contention window, which implicitly neglects collisions. Furthermore, synchronized stations are assumed that start their countdown procedures at the same time. The sums of uniform random variables are expressed using the Irwin-Hall distribution yielding the conditional probability that station one transmits k packets given station two transmits a single packet by

$$\mathsf{P}[K = k|1] = \frac{k+1}{(k+2)!}. \tag{2.3}$$

In [27] Berger-Sabbatel et al. reported a good match of the model with empirical data and conclude short-term fairness for two stations. The closed-form result, however, is not extended to incorporate inter-transmissions between more than two successive packets of a tagged station nor to more than two contending stations. In [28], the same authors provided measurement results that indicate a decreasing fairness for an increasing number of contending stations. They lack, however, an analytical model.

In this thesis we derive an extended version of the Berger-Sabbatel-metric to elaborate both short- and long-term fairness, which can hardly be done using the original version due to the limitations of the Irwin-Hall distribution, also in cases of more than two stations.

PROBLEM STATEMENT

Based on the literature review in Chapter 2, we extract several open research questions that are addressed throughout this thesis.

Regarding analytical performance analysis, we find some seminal results addressing QoS metrics, such as average throughput, medium access as well as one way delay, and fairness for two contending stations, achieved by the DCF [31, 41, 42, 96, 129]. In face of its success, however, it is remarkable that there is little consensus of the actual degree of fairness achieved, particularly bearing in mind its impact on quality of service [24, 27, 28, 64, 76, 97, 115, 141]. Until now, there is a lack of understanding on the parameters that influence fairness on long and short time scales. Moreover, to the best of our knowledge, there is no consistent model for larger networks, comprising more than two stations [27, 28]. To this end, we phrase the following questions.

> *How can the fairness of the DCF be described on any time scales as well*
> *as for a various number of contending nodes? What are the parameters*
> *that influence fairness? Is it possible to derive a fairness model that*
> *facilitates an application within the stochastic network calculus?*

The answers to those questions lead to a better understanding of the interaction of various wireless nodes. Since fairness has a tremendous impact on related QoS parameters, such as channel access times and one-way delays, this offers the possibility to improve network protocols that are prone to those effects, like TCP for example. We can also identify and tweak parameters, potentially online, to trade fairness to other metrics like throughput. Moreover, a stochastic network calculus model allows the analysis within the whole network calculus framework. To the end, we can derive end-to-end performance bounds and, therefore, pave the way to a holistic framework for network performance analysis.

Regarding service measurements in general and available bandwidth estimation in particular, we find some of the main research questions still unanswered. It has been shown before that infinite probing trains yield optimal results for stationary cross-traffic conditions [108, 109]. However, for networks with changing, potentially non-stationary, channel conditions, such as wireless networks, this might not be the case. To this end, we need shorter probing trains in oder to detect changes in the time-varying bandwidth process. The optimal train length, however, is still under discussion [72, 107]. Related to the question of packet train length is the question of convergence speed and the possibility of continuous probing. This also raises questions regarding the type of probing traffic and its intensity [72]. While Liebeherr et al. address the type of probing traffic in [107], the aspect of probing traffic intensity remains still open. However, this has a tremendous impact on the applicability of these methods, since we do not want to waste bandwidth unnecessarily. Bearing wireless networks in mind, we phrase the following questions.

> *What is the optimal probing method for service and available bandwidth estimation in wireless LANs that also allows for continuous measurements and online monitoring? What are the tuning knobs and influences of parameters on probing accuracy? How do we have to parameterize a probing technique, including its filters, to obtain optimal results?*

The answers to the above questions are relevant to several practical applications. A solid measurement methodology based on accurate theoretic fundamentals allows for an efficient QoS monitoring. To this end, applications such as measurement-based admission control and the verification of service-level agreements can benefit. Moreover, elastic traffic, like TCP, and elastic applications, like voice and video, can adapt their sending rate according to measurement results and, therefore, increase the user's quality of experience.

FAIRNESS ANALYSIS OF THE IEEE 802.11 DCF

In the following chapter, we evaluate the DCF that aims to achieve a fair and efficient medium access in IEEE 802.11. We provide an accurate model of the fairness of the DCF. Given M greedy stations we assume fairness if a tagged station contributes a share of $1/M$ to the overall number of packets transmitted. We identify and explain relevant characteristics and parameters of wireless links and the DCF, respectively, which have a severe influence on the fairness. We derive the probability distribution of fairness deviations. To this account, we extend the fairness model introduced by Berger-Sabbatel et al. [27] to accurately cover the system behavior of an IEEE 802.11 basic service set on long and short time scales. We use probability theory, see, e.g., [127], to derive closed-form expressions for the fairness that is achieved among M contending stations.

We tag one station $\hat{M}$, denote K_i the inter-transmissions of station $i = 1 \ldots M - 1$, and let $K = \sum_{i=1}^{M-1} K_i$. The conditional probability $P[K=k|l]$ that all other contending stations transmit k packets while the tagged station successfully transmits l packets, can be defined for $M \geq 2$ as

$$\mathsf{P}[K=k|l] = \mathsf{P}\left[\sum_{i=1}^{M-1} K_i = k \Big| l\right] \tag{4.1}$$

where the random variables K_i are the integers that satisfy

$$\sum_{j=1}^{K_i} b_i(j) \leq \sum_{j=1}^{l} b_{\hat{M}}(j) \quad \text{and} \quad \sum_{j=1}^{K_i+1} b_i(j) > \sum_{j=1}^{l} b_{\hat{M}}(j).$$

where $b_i(j)$ specifies the random backoff values of packet j on station i.

Applying our model, we first analyze long-term fairness with arbitrary backoff among two and more stations, i.e. $M \geq 2$ and $l, k \gg 1$. We compare backoff models with different backoff distributions and analyze their influ-

ence on fairness. Next, assuming exponentially distributed backoff values, we derive exact results for the distribution of inter-transmissions to cover short- and long-term fairness. Last, we calculate some valuable bounds and approximations that will be used by applications in Chapter 5 and Chapter 6. In addition, we support our analytical results by an extensive set of measurements.

4.1 RELEVANT WIRELESS LINK CHARACTERISTICS

In the following, we briefly discuss relevant characteristics of wireless links in general that are of vital importance for fairness as well as related applications such as available bandwidth estimation.

Fading and interference: Opposed to wired links, the characteristics of wireless channels are highly variable due to fading. Other, potentially hidden stations, which may include stations that implement different radio standards using the same frequency band, create interference on the wireless broadcast medium. These effects can cause rapid fluctuations of the signal-to-noise ratio and may lead to high bit error rates. Different modulation and coding schemes combined with rate adaptation may be used for compensation. As a consequence, the capacity and the availability of the channel may vary drastically. Furthermore, physical-layer capturing, that is implemented in real wireless devices, enables the correct reception of frames with a higher signal strength at the receiver, also in case of collisions. Evidently, this can have a tremendous effect on the performance of a wireless system, see, e.g., [65, 65] and references therein.

Contention: In case of wireless multi-access channels, stations share the same medium and contend for access to the channel controlled by the MAC protocol. Before accessing the medium stations listen to the channel to detect nearby transmissions with the objective of avoiding collisions. This procedure may fail in case of hidden stations, thus requiring additional protocol mechanisms such as RTS/CTS. The resulting behavior of medium

access procedures may be largely different if compared to FCFS multiplexing at a point-to-point link.

Retransmissions: Due to frequent packet loss on wireless links, e.g., because of fading, interference, or collisions, the IEEE 802.11 standard includes an RLC protocol that implements an automatic repeat request technique, i.e. stop-and-wait ARQ, to ensure packet delivery. Link layer retransmissions consume channel capacity and lead to increased and varying one-way delays.

Our highly controlled measurement environment for IEEE 802.11 networks, however, eliminates effects that are due to fading, interference from external sources, and physical layer capturing to a large extent. Protocol-related aspects, however, remain to be addressed.

4.2 EMPIRICAL FAIRNESS EVALUATION

As mentioned before, the fair capacity allocation of IEEE 802.11 links using the DCF is discussed controversially in the literature. While [93] showed short-term unfairness of CSMA/CA-based WaveLANs a recent study [27] attributes findings of unfairness to early WaveLAN cards and reports that current IEEE 802.11 DCF implementations actually exhibit good short-term fairness. On this account, we conduct an empirical evaluation of the fairness achieved by the DCF that acts as a baseline for our model-based analysis. We perform extensive experiments with two and more contending stations using an IEEE 802.11 testbed, partly in a shielded unechoic measurement chamber. For comparison, we present results for the same topology from simulations with the OMNeT++ [7] network simulator[1]. Using a very simple simulation model for the PHY layer, this allows the analysis of wireless LANs with a perfect physical channel.

---

1 We used OMNeT++ version 4.1 and the IEEE 802.11 simulation models that are part of the current INET-MANET [4] framework for wireless network simulations. These models are based on the IEEE 802.11g extension from Cocorada [48].
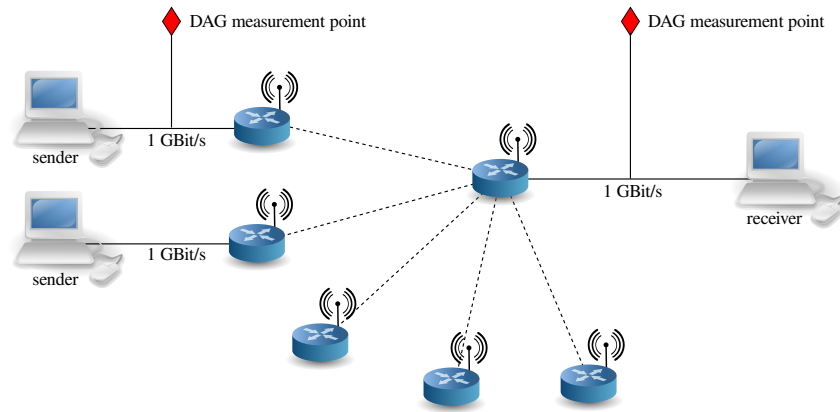
Figure 6: Wireless testbed for empirical fairness evaluation comprising one access point and up to 10 wireless clients connected using the IEEE 802.11a standard achieving an overall maximum goodput of approx. 28 Mbps.

Our testbed depicted in Figure 6 consists of up to 10 wireless stations that ran the iperf traffic generator [5] to send data to a receiver and contend for the medium. The stations[2] are connected to the access point using IEEE 802.11a with a fixed rate of 54 Mbps. In addition, the channel is monitored by a spectrum analyzer to detect interferences that may arise from outside the testbed. The access point is connected to the receiver using fast Ethernet with 1 Gbps. The distance between the wireless stations and the access point is between 0.5 m and 1.5 m to ensure that all stations are in radio range to each other and to avoid antenna near-field effects. Furthermore, we adapted the sending power such that the signal strength of all stations is equal on average at the receiver. We note that for an increasing number of stations this correct calibration of the testbed becomes absolutely relevant to reduce physical-layer influences. We switched off RTS/CTS, automatic rate adaption, and made sure that packet fragmentation does not occur. Additionally, all stations are connected to a separated wired control network and we used SSH-Launcher [35] scripts to automate our experiments. In addition, we applied similar simulations comprising up to 40 wireless nodes.

---

2 We used Gateworks Laguna GW 2388-4 Boards [3] with a 600 MHz Dual Core ARM11 processor and 128 MB RAM running OpenWrt [8] Kamikaze bleeding edge at revision 27943 with Linux kernel version 2.6.39 [6]. We employed Ubiquiti SR-71A 802.11abgn wireless LAN cards.

(a) two contending flows

(b) four contending flows

Figure 7: Measured throughput of contending flows over 60 s. Each flow obtains a long-term fair bandwidth share. Regarding the simulations, we attain the very same results.

Unless otherwise noted, we carried out all experiments lasting for 15 minutes per participating station. We repeated each experiment 25 times to generate a sufficient amount of data for statistical analysis. Each station sent a greedy UDP flow of 1500 Byte packets at a rate of 28 Mbps. Note that this sending rate of a single station coincides with the service rate of IEEE 802.11g due to per-packet protocol overhead.

Figure 7 shows the long-term average throughput of contending flows with a constant bit rate. The throughput is averaged over 25 experiments that last for 60 s each. We also calculated confidence intervals at a confidence level of 0.95. Since these intervals are negligibly small, they are, however, omitted in the plots. In Figure 7a two flows contend for the link. Flow 1 has a rate of 28 Mbps and the rate of flow 2 increases from 0 to 28 Mbps in steps of 1 Mbps for each experiment. Similarly, Figure 7b shows the throughput of four flows, where the rate of flow 4 increases the same way.

The results in Figure 7 confirm that each flow achieves a fair share of the capacity on long time scales. Flow 2 in Figure 7a achieves its target throughput whereas the throughput of flow 1 is reduced accordingly until flow 2 reaches 14 Mbps. From this point on both flows get a fair share of 14 Mbps regardless of the rate of flow 2. Figure 7b confirms this result for four heterogeneous flows. We note that throughout all experiments and simulations, the stations transmitted their data with an average rate

(a) pmf of inter-transmissions
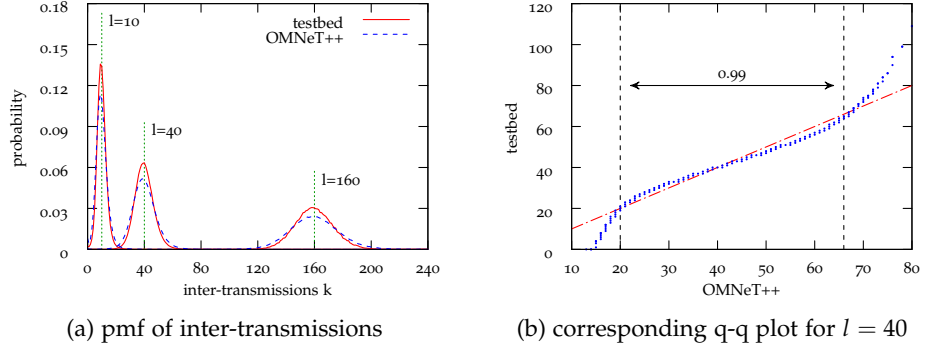
(b) corresponding q-q plot for $l = 40$

Figure 8: OMNeT++ simulation results agree closely with the measurement data. The q-q plot matches well for 0.99 of the samples, but brings out deviations at the tail, showing additional unfairness in the testbed compared to the simulator. Moreover, the slope midway does not become one, which indicates a better fairness in the testbed.

of approximately $28/M$ Mbps and a total deviation from that of less than $\pm 2.5\%$. This indicates close to perfect long-term fairness in all cases.

To investigate fairness at different timescales, we tag one station and count the random number of packet inter-transmissions K of all other contending stations while the tagged station transmits l packets. To this end, we consider the conditional probability $P[K = k|l]$ where mean and variance are related to the fairness. That is, the smaller the variance of the inter-transmissions K and the more similar the mean of K becomes to l, the better the fairness. Hence, $E[K] = l$ and $Var[K] = 0$ on all time scales, i.e. for all l, indicate perfect fairness. For $M = 2$ and $l = 1$ our definition reduces to the special case in [27] where short-term fairness among two stations is analyzed. Throughout this thesis, however, we consider short- as well as long-term fairness and an arbitrary number of stations.

In the following, we compare our testbed measurements to OMNeT++ simulations. Figure 8a shows the probability mass function (pmf) of the inter-transmissions K for different l. We observe that the simulation results match the measurement data quite well but underestimate the fairness slightly. We use quantile-quantile (q-q) plots to bring out differences at the tail of the distributions to detail the goodness of fit. If the q-q plot equals a straight line, the distributions are similar in general, however, they may

(a) pmf of inter-transmissions
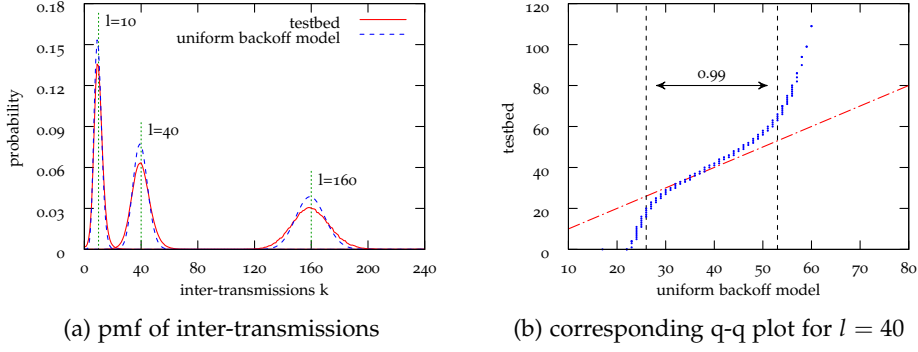
(b) corresponding q-q plot for $l = 40$

Figure 9: Numerical results from an extension of the model proposed by Berger-Sabbatel et al. in [27]. The uniform backoff model overestimates fairness. The q-q plot not only differs at the tail but also its slope midway does not become one. Moreover, the 0.99 intervall narrows by a factor of 1.7 and wrongly indicates good fairness

have different distribution parameters. The slope of the line is related to the variance whereas its position reflects the mean. Thus, a straight line with slop 1 and $x = y$ represents equal distributions with identical parameters. The exemplary q-q plot for $l = 40$ in Figure 8b shows the quantiles of the measurement data vs. the simulation data. As indicated in Figure 8b we find that 0.99 of the samples almost coincide. The slightly different slope confirms the underestimation of fairness stated before. Deviations at the tail, however, show additional unfairness in the testbed that is not reproduced by the simulator, i.e. in the testbed large deviations of K occur more frequently.

Furthermore, we compare our testbed results to the model that is established for two stations and short-term fairness in [27] and discussed in Chapter 2. To analyze long-term fairness as well we extend this method from the special case $l = 1$ in [27] to cover all $l \geq 1$. Denote $b_i(j)$ the i.i.d. countdown values for packet j at station $i = 1, 2$. We model the distribution of inter-transmissions K of station 1 while station 2 transmits l packets as

$$P[K = k|l] = P\left[\sum_{j=1}^{k} b_1(j) \leq \sum_{j=1}^{l} b_2(j) \text{ and } \sum_{j=1}^{k+1} b_1(j) > \sum_{j=1}^{l} b_2(j)\right]. \tag{4.2}$$

Assuming $b_i(j)$ are uniform random variables as in [27], equation (4.2) results in Irwin-Hall distributed terms which, however, do not yield a simple solution for $l \geq 2$. For now, we compute the distribution numerically by convolution and compare the results denoted *uniform backoff model* to our testbed data in Figure 9. The pmf of the model shows a deviation from the testbed results that is also confirmed by the q-q plot. Compared to Figure 8b the 0.99 interval becomes much narrower in Figure 9b indicating that the assumption of uniform backoff values overestimates the actual fairness of the DCF.

4.3   LONG-TERM FAIRNESS WITH ARBITRARY BACKOFF

We use the central limit theorem, that states that any infinite sum of i.i.d. random variables becomes normally distributed, to derive the long-term fairness, that is related to the average throughput, with arbitrary backoff. In the sequel, we denote normal random variables $N(\mu, \sigma^2)$ where μ is the mean and σ^2 the variance.

Theorem 1 (Gaussian approximation) *Let $b_i(j)$ be i.i.d. random variables with mean μ and variance σ^2 and let $M = 2$. For $k, l \gg 1$ the cumulative distribution function of Eq. (4.1) is approximately Gaussian with*

$$P[K \leq k|l] \approx P\left[N(0,1) \leq \frac{\mu\,(k-l)}{\sigma\,\sqrt{k+l}}\right].$$

Proof   For $M = 2$ we have from equation (4.1) that

$$P[K < k|l] = P\left[\sum_{j=1}^{k} b_1(j) > \sum_{j=1}^{l} b_2(j)\right]$$

and after expanding and normalizing this formula equals

$$= P\left[\frac{\sum_{j=1}^{l} b_2(j) - l\mu}{\sigma\sqrt{l}} - \frac{\sum_{j=1}^{k} b_1(j) - k\mu}{\sigma\sqrt{l}} < \frac{\mu(k-l)}{\sigma\sqrt{l}}\right].$$

Applying the central limit theorem the sums can be approximated by normal distributions with specific means and variances. Thus, it follows that

$$P[K < k|l] \approx P\left[N(0,1) - N\left(0,\frac{k}{l}\right) < \frac{\mu(k-l)}{\sigma\sqrt{l}}\right].$$

Since the normal distribution with zero mean is symmetric we can replace the subtraction of $N(0,k/l)$ by an addition. Furthermore, the sum of two normal random variables $N(\mu_1, \sigma_1^2)$ and $N(\mu_2, \sigma_2^2)$ is also normal with $N(\mu_1 + \mu_2, \sigma_1^2 + \sigma_2^2)$ such that

$$P[K < k|l] \approx P\left[N\left(0,\frac{k+l}{l}\right) < \frac{\mu(k-l)}{\sigma\sqrt{l}}\right].$$

Finally, we use that if X is $N(a\mu, a^2\sigma^2)$ then $Y = X/a$ is $N(\mu, \sigma^2)$ with $a^2 = (k+l)/l$ to standardize the result. ∎

It is worth highlighting that Theorem 1 assumes i.i.d. random countdown values. It does, however, not make any assumption about their distribution. To compare the impact of different backoff models on fairness we phrase the following corollary for uniformly, as used, e.g., in the short-term fairness model in [27], and exponentially distributed countdown values, as assumed, e.g., by the throughput models in [31, 95] respectively.

Corollary 1 (Uniform versus exponential countdown) *Assume Theorem 1. If the $b_i(j)$ are uniform in $[0, w]$, i.e. $\mu = w/2$ and $\sigma^2 = w^2/12$, then*

$$P[K \leq k|l] \approx P\left[N(0,1) \leq \frac{\sqrt{3}(k-l)}{\sqrt{k+l}}\right].$$

If the $b_i(j)$ are exponentially distributed with parameter λ, i.e. $\mu = \lambda^{-1}$ and $\sigma^2 = \lambda^{-2}$, then

$$P[K \leq k|l] \approx P\left[N(0,1) \leq \frac{k-l}{\sqrt{k+l}}\right].$$

Theorem 1 and Corollary 1 yield a number of important conclusions. First, we compare the pmfs from Corollary 1 displayed in Figure 10a. In case of two stations, the assumption of exponential backoff values, which reflects the increasing contention window and retransmissions in case of collisions, matches the simulation results perfectly. For the testbed, however, it underestimates the fairness and yields a rather pessimistic result. In contrast the assumption of uniform backoff overestimates the fairness. These results are owing to the fact that in case of two contending stations only, the number of collisions and retransmissions is quite low. Thus, the DCF's backoff values are sampled from a uniform distribution in the interval $[0, w_{min} - 1]$, i.e. without an increasing contention window, in most cases. Hence, for just two stations the distribution of inter-transmissions lies in between the uniform and exponential backoff approximation. For more than two stations, however, the number of collisions and retransmissions increases and the binary exponential backoff mechanism comes into account. In this case, the exponential backoff model provides interesting results and insights. This is further elaborated in Section 4.4. The uniform backoff model, however, cannot be transformed into a closed form solution for more than two hosts. Hence, it does not provide any further results.

Regarding the goodness of fit of the exponential backoff approximation for two stations, Figure 10c and Figure 10d show q-q plots of the exponential model vs. the testbed data and confirm the accuracy of the model for 0.99 of the samples. The testbed exhibits larger unfairness at the distribution tail as also observed for the uniform model compared to OMNeT++ simulations in Figure 8.

Next, we consider Corollary 1 and find that the distribution parameter in case of uniform as well as in case of exponential countdown values has no influence on the fairness. In contrast, the distribution itself has significant impact. Corollary 1 shows an explicit fairness degradation of $\sqrt{3}$ of exponential compared to uniform countdown values, i.e. $\sqrt{3}$ can be viewed as the price of exponential backoff and retransmissions. Figure 10a

(a) pmf of inter-transmissions yield a close fit for the exponential backoff approximation

(b) q-q plot for $l = 40$ vs. $l = 160$ showing the fairness improvement with $\sqrt{l}$

(c) corresponding q-q plot for $l = 40$

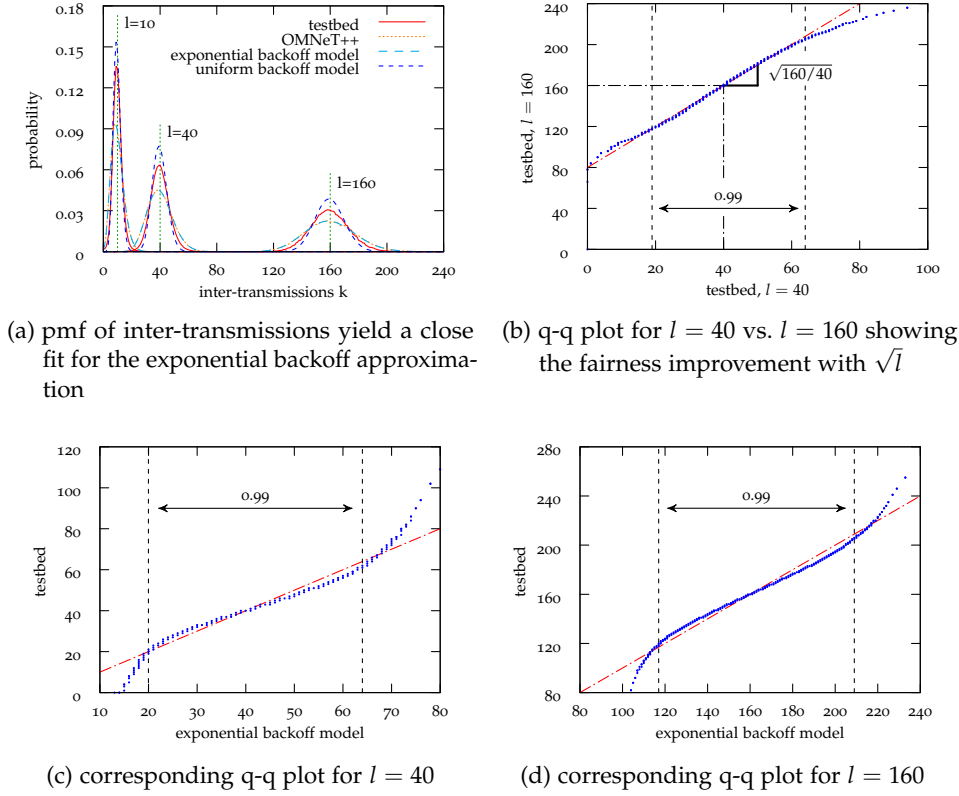(d) corresponding q-q plot for $l = 160$

Figure 10: Using exponential backoff values, the analytical results from Corollary 1 match the OMNeT++ results perfectly and the testbed data quite well. Assuming uniformly distributed random backoff variables the model overestimates the fairness and yields an overoptimistic result. Furthermore, the model accurately predicts the $\sqrt{l}$ fairness improvement, as illustrated by the slope in Figure 10b.

shows this effect clearly, i.e. the pmf for uniform is by $\sqrt{3}$ higher and narrower.

Last, we evaluate the improvement of long-term over short-term fairness. To this end, we define a multiplicative constant c and let $k = cl$. The parameter c may be viewed as a threshold value that specifies a relative deviation that is still considered fair. By insertion the term $(k - l)/\sqrt{k + l}$ from Theorem 1 becomes $\sqrt{l}\,(c - 1)/\sqrt{c + 1}$ and it follows that long-term fairness improves proportionally to $\sqrt{l}$. Thus, the initial fairness improvement for small l is significant but becomes less pronounced with increasing l. This result is independent of the distribution of backoff values. Figure 10b depicts a q-q plot of our measurement data for $l = 40$ vs. $l = 160$. The slope of the

q-q plot closely follows $\sqrt{160/40}$ for 0.99 of the samples. Hence, Figure 10b clearly shows the $\sqrt{l}$ scaling in the testbed data.

4.4   SHORT- AND LONG-TERM FAIRNESS WITH EXPONENTIAL BACKOFF

In the sequel we consider the exponential backoff model that is widely used in the literature, e.g., in [31, 41] for throughput analysis. Moreover, it has been proven fairly accurate for the Gaussian approximation to our empirical results in Section 4.3. Opposed to the uniform backoff model, the exponential backoff model allows an analytical solution of equation (4.1) for $M > 2$ that can be further elaborated. To this end, we derive an exact result and useful approximations for long- and short-term fairness respectively.

Theorem 2 (Exact result) *Let $b_i(j)$ be i.i.d. exponential random variables and let $p = 1/M$. Then Eq. (4.1) is negative binomial*

$$P[K = k|l] = p^l (1 - p)^k \binom{k + l - 1}{k}.$$

Proof   First we provide a direct proof for two contenting stations, i.e. $M = 2$, and relax this assumption by using an intuitive argumentation in the latter. The direct proof follows readily by using equation (4.1)

$$P[K = k|l] = P\left[\sum_{j=1}^{k} b_1(j) \leq \sum_{j=1}^{l} b_2(j) \text{ and } \sum_{j=1}^{k+1} b_1(j) > \sum_{j=1}^{l} b_2(j)\right]$$

that can be rewritten as

$$P[K = k|l] = P[0 \leq Z - Y < X].$$

Here, X, Y, Z are independent random variables where $X = b_1(k+1)$ is exponentially distributed with a probability density

$$f_X(x) = \lambda e^{-\lambda x}$$

and $Y = \sum_{j=1}^{k} b_1(j)$ as well as $Z = \sum_{j=1}^{l} b_2(j)$ are sums of k and l exponentially distributed random variables respectively. Hence, Y and Z are Gamma (also known as m-Erlang) distributed with the well-known probability density functions given by

$$f_Y(y) = \frac{\lambda e^{-\lambda y} (\lambda y)^{k-1}}{(k-1)!}$$

$$f_Z(z) = \frac{\lambda e^{-\lambda z} (\lambda z)^{l-1}}{(l-1)!}.$$

We calculate the probability $P[K=k]$ as follows

$$P[K=k|l] = \iiint_{0 \leq z - y < x} f_Z(z) f_Y(y) f_X(x) dx dy dz$$

$$= \int_{z=0}^{\infty} \int_{y=0}^{z} \int_{x=z-y}^{\infty} f_Z(z) f_Y(y) f_X(x) dx dy dz$$

Solving the first two integrals yields

$$P[K=k|l] = \frac{\lambda^{k+l}}{(l-1)! \, k!} \int_{z=0}^{\infty} e^{-2\lambda z} z^{k+l-1} \, dz.$$

We solve this equation using integration by parts, and after the first step this yields

$$P[K=k|l] = \frac{\lambda^{k+l}}{(l-1)! \, k!} \frac{(k+l-1)}{2\lambda} \int_{z=0}^{\infty} e^{-2\lambda z} z^{k+l-2} \, dz$$

which can be further simplified by continuously applying the integration by parts $k + l - 2$ times. Finally this yields

$$P[K=k|l] = \frac{\lambda^{k+l}}{(l-1)! \, k!} \frac{(k+l-1)!}{(2\lambda)^{k+l-1}} \int_{z=0}^{\infty} e^{-2\lambda z} \, dz$$

where the last integral evolves to $1/2\lambda$ such that

$$P[K=k|l] = \frac{1}{(l-1)! \, k!} \frac{(k+l-1)!}{2^{k+l}}$$

and after some reordering

$$P[K=k|l] = 2^{-(k+l)} \binom{k+l-1}{k}$$

becomes negative binomial.

Next, we relax the assumption of considering only two contending stations and provide an intuitive proof that Theorem 2 holds for $M \geq 2$ in general. Therefore, we determine the probability that station $\hat{M}$ gets access to the channel. Key to derive this probability is the memoryless property of negative exponential random variables, i.e. given an exponential random variable X with parameter λ it holds that $P[X > x + y | X > x] = P[X > y] = e^{-\lambda y}$, see, e.g., [127].

Consider M stations that contend for an idle channel. Owing to the memoryless property, each station has an exponentially distributed countdown value with the same parameter λ irrespective of the time the station has already spent on performing the countdown procedure. It follows that each channel access can be viewed as an independent Bernoulli experiment. Denote p the probability of success, i.e. the probability that station $\hat{M}$ finishes its countdown procedure first, such that it attains access to the channel. Since the remaining countdown values are i.i.d. at all stations, each station has the same channel access probability p where $\sum_{i=1}^{M} p = 1$ such that $p = 1/M$. Hence, the probability that stations $1 \ldots M - 1$ access the channel exactly k times until station $\hat{M}$ performs the l-th channel access is the probability of seeing the l-th success of station $\hat{M}$ exactly in the $(k+l)$-th Bernoulli trial. This event is negative binomially distributed. ∎

Figure 11 compares the results from Theorem 2 with the testbed measurement data. We find that Theorem 2 predicts short- and long-term fairness for $M = 2$ stations fairly accurate, see Figure 11a. Figure 11c and Figure 11d add q-q plots for short-term fairness which also show a close match for 0.99 of the samples. Figure 11b restates the $\sqrt{l}$-effect also for short time scales.

(a) pmf of inter-transmissions, $M = 2$

(b) corresponding q-q plot for $l = 1$ vs. $l = 10$

(c) corresponding q-q plot for $l = 1$

(d) corresponding q-q plot for $l = 10$

Figure 11: Analytical results from the exponential model and Theorem 2 predict short- and long-term fairness correctly for $M = 2$ stations. As already stated before, fairness improves with $\sqrt{l}$.

Next, we derive a normal approximation. This allows viewing fairness deviations as i.i.d. Gaussian noise.

Corollary 2 (Gaussian approximation) *Assume Theorem 2. For $l \gg 1$ it follows that*

$$\mathsf{P}[K \leq k | l] \approx \mathsf{P}\left[N(0,1) \leq \frac{kp - l(1-p)}{\sqrt{l(1-p)}}\right].$$

Proof We view the negative binomial random variable in Theorem 2 as a sum of i.i.d. geometric random variables. Each of the geometric random variables equals the number of trials required until the next success is achieved, i.e. we write Theorem 2 as a sum of l i.i.d. geometric random variables denoted X_i

$$\mathsf{P}[K \leq k | l] = \mathsf{P}\left[\sum_{i=1}^{l} X_i \leq k + l\right].$$

(a) Kullback-Leibler distance for approxi-
mations vs. exact result

(b) Kullback-Leibler distance for analytical
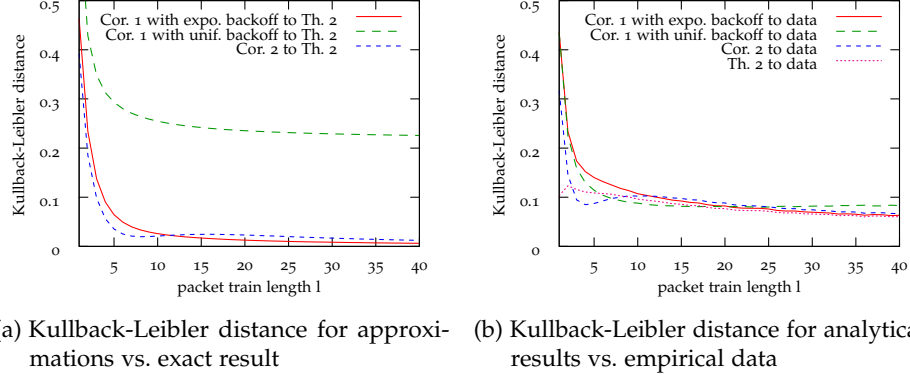results vs. empirical data

Figure 12: The Kullback-Leibler distances quantify the goodness of fit of the ap-
proximations and analytical results to our testbed data. We find that the
Gaussian approximations assuming exponentially distributed backoff
values match the analytical results as well as the testbed very closely.
Moreover, the same holds for the exact results, which is in line with the
q-q plots shown before. We conclude that we can apply the Gaussian
approximations already for small packet train length, i.e. $l > 4$.

Normalization using the mean $\mu = 1/p$ and the variance $\sigma^2 = (1-p)/p^2$
of geometric random variables yields

$$\mathsf{P}[K \leq k | l] = \mathsf{P}\left[\frac{\sum_{i=1}^{l} X_i - l\mu}{\sigma\sqrt{l}} \leq \frac{(k+l)p - l}{\sqrt{l(1-p)}}\right].$$

Using the central limit theorem the normalized sum is approximately
standard normal $N(0,1)$ if $l \gg 1$. ∎

We note that Corollary 1 and Corollary 2 for $M = 2$ converge under
the assumption of the central limit theorem, i.e. for $k, l \rightarrow \infty$. For further
quantitative comparison we use the Kullback-Leibler distance [49] between
the measurement data and the analytical expressions. The Kullback-Leibler
distance $D(X||Y)$ quantifies the deficiency if we assume the distribution of
X instead of the true distribution of Y. It is defined as

$$\mathsf{D}(X||Y) = \sum_{x} \mathsf{P}(X = x) \ln \frac{\mathsf{P}(X = x)}{\mathsf{P}(Y = x)}.$$

Here, $P(X = x)$ is either the empirical mass function or the negative binomial
distribution, i.e. the exact results, and $P(Y = y)$ is the probability mass
function from the analytical model.

Figure 12 summarizes the results. Figure 12a compares the Gaussian approximations of Corollary 1 and Corollary 2 to the exact results of Theorem 2. We find that the Kullback-Leibler distance converges to zero very fast. Hence, the Gaussian approximations that assume exponentially distributed backoff values match the exact results already for medium-sized packet trains. We conclude that we can apply the Gaussian approximations for medium packet train length, i.e. $l > 4$. Similar to the q-q plots presented above, Figure 12b illustrates the goodness of fit of the analytical results to the testbed data. Again, the exact result matches the testbed data very well.

Finally, we derive a useful bound for the distribution of inter-transmissions. In the following chapter, this bound is used to compute a stochastic service curve that describes the service offered by the DCF.

Corollary 3 (Chernoff bound) *Assume Th. 2. It follows that*

$$\mathsf{P}\left[K \lesseqqgtr k \big| l\right] \leq \left(\frac{(1-p)(k+l)}{k}\right)^k \left(\frac{p(k+l)}{l}\right)^l \quad \forall k \lesseqqgtr l(M-1)$$

Proof The proof uses Chernoff bounds, see, e.g., [127]

$$\mathsf{P}\left[X \lesseqqgtr x\right] \leq e^{-\theta x} \mathsf{M}_X(\theta) \quad \forall \theta \lesseqqgtr 0$$

that provide exponentially decreasing bounds on tail distributions and yields good results for small probabilities. Here, $\mathsf{M}_X(\theta) = E[e^{\theta X}]$ denotes the moment generating function of X. We insert the well-known moment generating function of the negative binomial distribution and derive

$$\mathsf{P}[K \leq k | l] \leq e^{-\theta k}\left(\frac{p}{1-(1-p)e^{\theta}}\right)^l \quad \forall \theta < 0.$$

In order to obtain the best-possible bound, we minimize the right-hand side over all $\theta < 0$ and insert $\theta = \ln(k) - \ln((l+k)(1-p))$ where $k < l(1-p)/p = l(M-1)$ to ensure $\theta < 0$. The upper bound follows in the same way. ∎

(a) Jain's index over packet train length

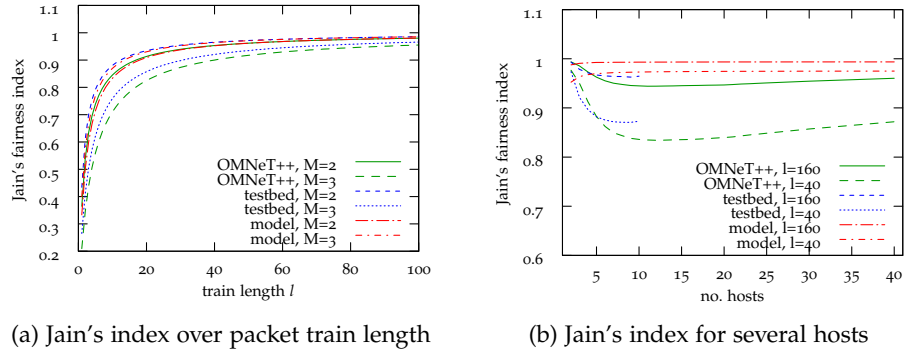(b) Jain's index for several hosts

Figure 13: Jain's fairness index over packet train length for various number of contending hosts. We find that fairness increases for an increasing number of packets per train, i.e. for long time scales. We note that this fairness improvement is achieved quite fast. Moreover, we find that fairness initially decreases for an increasing number of hosts. At a first glance, this seems to be contrary to the model outcome. However, the fairness increases again, for a larger number of hosts as correctly predicted by the model.

Theorem 2, Corollary 2, and Corollary 3 can be easily extended to heterogenous stations that use different parameters λ, e.g., for service differentiation. In this case only the probability of successful channel access p has to be adapted accordingly. Also, it is a straightforward extension of Theorem 2 to derive the probability that a single station with index M transmits l packets given that the remaining $M-1$ stations together transmit k packets.

To further elaborate these fairness issues, also in case of more than two stations, we derive Jain's fairness index, see equation (2.2), which follows directly from the first and second moment of the inter-transmissions. For the negative binomial distribution in Theorem 2 the first moment is $E[K] = l(1-p)/p$ and the second central moment is $l(1-p)/p^2$ such that the second moment becomes $E[K^2] = l(1-p)/p^2 + (l(1-p)/p)^2$. With $p = 1/M$ we have

$$J = \frac{l}{l + \frac{M}{M-1}}. \tag{4.3}$$

Regarding Figure 13, we find that equation (4.3) matches the testbed measurement data almost perfectly for $M = 2$ stations, in which case

$J = l/(l + 2)$. In face of short-term unfairness we find that fairness is practically achieved already for moderate train length l, see Figure 13a. For more than two stations additional effects that cause unfairness result in a deviation. However, as depicted in Figure 13b, this deviance diminishes and the system becomes fairer again for an increasing number of stations as correctly predicted by the model. We conclude that additional effects beyond our model cause unfairness in cases of more than two stations. Empirical results indicating poor fairness in case of more than two stations have also been reported, e.g., in [28]. We assume, however, that theses effects become negligible again for $M \gg 2$.

To elaborate the fairness deviation in case of $2 < M < \infty$ stations, we recall the assumptions of the exponential backoff model. As already stated by the decoupling approximation introduced by Bianchi [31] and Cali [41] respectively, the exponential backoff model assumes i.i.d backoff values that are implicitly memoryless which also results in independent packet arrivals. In real systems, however, the backoff processes of various stations are coupled by collisions and are not memoryless. For example, consider a network with two stations that sample their backoff values from a uniform distribution in the range of $[0, w_{min} - 1]$ and proceed with the standard medium access procedure. In case one station successfully transmitted its packet, it is more likely that the other stations transmits next, since the system is not memoryless. Moreover, if we consider three stations, we find that if the packets of two stations collide, these stations double their contention window and the third station can take advantage from this. Thus, the stations are clearly coupled by their collisions. We conclude that packet arrivals in real wireless systems are not independent but correlated, which violates the model assumptions to some extent.

To analyze and quantify these correlations we define a packet arrival process $\kappa(n)$ for $n = 1, 2, ..., k$ where $\kappa(n) = 1$ if a tagged station sends a

(a) Autocorrelations of packet arrivals over lag ι measured in the testbed
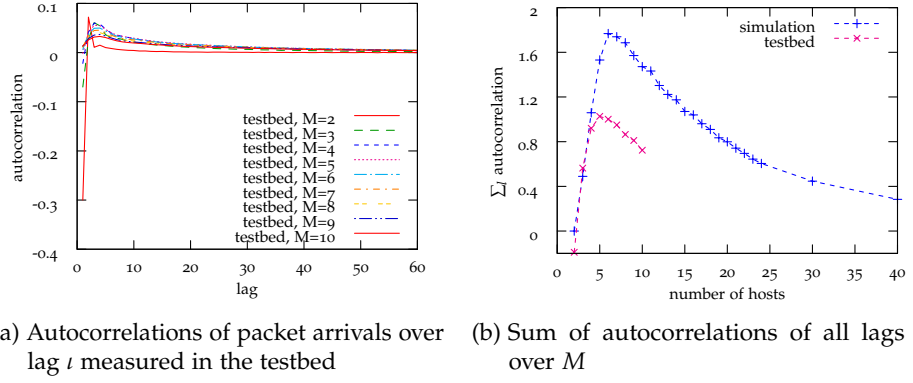
(b) Sum of autocorrelations of all lags ι over M

Figure 14: Autocorrelation of packet arrivals over lag and the sum of autocorrelation over number of hosts from measurements and simulations respectively. We find that the correlation of packet arrivals decreases if the lag increases. Thus, the packet arrival process becomes uncorrelated after a number packets are sent. The sum of autocorrelations remains finite and decreases again for a larger number of hosts. Hence, the model assumptions are met for $M \gg 2$ which explains the minor model inaccuracies for a medium number of stations.

packet and $\kappa(n) = -1$ if any other station sends a packet. Given the process is stationary, we calculate the sample autocorrelation $\hat{R}(\iota)$ as [99]

$$\hat{R}(\iota) = \frac{1}{Var(\kappa)\,(k-\iota)} \sum_{n=1}^{k-\iota} [\kappa(n) - \overline{\kappa}]\,[(\kappa(n+\iota) - \overline{\kappa}]$$

where $\overline{\kappa}$ and $Var(\kappa)$ are the sample mean and variance of $\kappa(n)$ respectively. The positive variable $\iota = 1, 2, ..., k-1$ denotes the lag that describes the time shifting of the process to itself. If, on the one hand, the observations of $\kappa(n)$ are independent, $\hat{R}(\iota)$ becomes zero. If, on the other hand, $\hat{R}(\iota)$ deviates from zero significantly, this is a strong indication that $\kappa(n)$ is not independent. Note, however, that the sample autocorrelation $\hat{R}(\iota)$ must not be exactly zero even if $\kappa(n)$ is independent due to the limited number of observations.

Figure 14a depicts the mean sample autocorrelation obtained by measurements of $M = 2, ..., 10$ hosts. We find that for all cases, i.e. for all setups with respect to M, $\kappa(n)$ is an *a-dependent process* [119] since $\hat{R}(\iota)$ becomes approximately zero after a steps. Thus, the packet arrival process becomes uncorrelated after a packets are sent. Furthermore, Figure 14b shows the

(a) q-q plot for $l = 40$ vs. $l = 160$, $M = 3$

(b) q-q plot for $l = 40$ vs. $l = 160$, $M = 4$

(c) q-q plot for $l = 40$ vs. $l = 160$, $M = 5$

(d) q-q plot for $l = 40$ vs. $l = 160$, $M = 10$

Figure 15: Long-term fairness improves approximately with $\sqrt{l}$ also in case of more than two stations. This improvement is correctly predicted by Corollary 2.

sum of autocorrelations over all ι, i.e. $\sum_{\iota=1}^{\infty} \hat{R}(\iota)$ for the measurements and the simulations respectively. In case this sum tends to infinity, it would have a tremendous impact on the statistical inferences, see, e.g., [25]. However, we see from Figure 14b that $\sum_{\iota=1}^{\infty} \hat{R}(\iota) = constant < \infty$ and, in addition, decays to zero for an increasing number of nodes after it has exceeded a certain maximum value. This indicates that for $M \gg 2$ the real system behavior converges to the model assumptions. For $M < \infty$ stations, we conclude that exponentially distributed backoff values model the real backoff process quite, but not perfectly accurately. This explains the slight fairness deviations of the testbed data compared to the model outcome for a medium number of stations. It is, however, worth highlighting that the analytical results keep hold of their validity and shed light on interesting insights regarding the DCF, also for a moderate number of stations.

Considering Theorem 2 we recover the result that the parameter of the i.i.d. exponentially distributed countdown values does not impact fairness

regardless of the number of stations and the time-scale. To this end, we again define a multiplicative constant c and let and $k = (M-1)cl$ using that $p = 1/M$. Thus, Corollary 2 yields

$$P[K \leq k|l] \approx P\left[N(0,1) \leq \sqrt{l}\,(c-1)\,\sqrt{\frac{M-1}{M}}\right].$$

predicting that fairness improves with $\sqrt{l}$. Figure 15 shows the improvement for $M = 3$, $M = 4$, $M = 5$, and $M = 10$ stations respectively. The testbed measurements confirm the dependence on $\sqrt{l}$.

DCF CLOCK AND A SERVICE CURVE MODEL

In this chapter we introduce the basic principles of deterministic and stochastic network calculus. Equipped with the findings of the previous chapter, we derive a service curve model for the DCF that facilitates applications of the stochastic network calculus [47, 58, 79]. To this end, we view the DCF as emulating the Generalized Processor Sharing discipline. Subsequently, we underpin our analytical results with an extensive set of measurements.

5.1 BACKGROUND ON NETWORK CALCULUS

The pioneering GPS model [120] defines a weighted resource allocation that is perfectly fair on any time scale. To this end, GPS assumes infinitesimal small packets, also referred to as fluid traffic, which can be divided, stopped, and interrupted at any time. Obviously, these characteristics cannot be found in real systems since traffic consists of packets with a certain length that occupy the complete channel for a period related to that length. To date, a variety of packet-by-packet implementations exist that emulate GPS closely, such as Weighted Fair Queuing (WFQ) [53, 120] and Deficit Round Robin (DRR) [134]. Distributed emulations are proposed in [23, 97, 141].

GPS and the calculus for network delays [51] gave rise to the important concept of service curves [120] that is the foundation of today's network calculus. Here, the notion of service curves characterizes the behavior of a time-invariant min-plus linear systems. The particular strength of network calculus is the convolution of tandem systems that yields the notion of network service curve and permits analyzing entire networks as a single system. Obviously, it is closely related to classical linear systems theory where the departing signal can be calculated by convolving the input signal with the impulse response of the system, see, e.g., [63, 123] for

details. This analogy is also illustrated, e.g., in [46, 59, 107]. For a detailed introduction to deterministic and stochastic network calculus we refer to the textbooks [43, 80, 101] as well as to [46, 59] and references therein.

In the following we explore an intuitive example of a so-called latency-rate service curve of a constant rate link. Therefore, the functions $A(\tau, t) = A(t) - A(\tau)$ and $D(\tau, t) = D(t) - D(\tau)$ denote the cumulative arrivals and the cumulative departures in bits of a system in the interval $[\tau, t)$. For simplicity we write $A(t)$ and $D(t)$ meaning $A(0, t)$ and $D(0, t)$. We assume that $A(t)$ and $D(t)$ are non-negative and monotonically increasing functions and from causality it reveals that $A(t) \geq D(t)$ for all t. The backlog, i.e. all bits that reside within the system at a given point in time, can be easily calculated by $B(t) = A(t) - D(t)$. The modeling time can be either continuous or discrete, where a continuous time arrival or departure function can be mapped to its discrete counterpart by sampling. This, however, results in a loss of information. Given a discrete time function, a continuous time signal can be reconstructed as well.

Considering a work-conserving unlimited-buffered constant rate link, which forwards fluid arrivals with a specific service rate r immediately whenever the system is backlogged, i.e. when there is data available, the departures leave the system with a rate that equals the service rate. Hence, for any two time instances $t \geq \tau \geq 0$ in any backlog interval $[\tau, t)$ it holds that

$$D(t) \geq D(\tau) + r(t - \tau) \tag{5.1}$$

stating that the cumulative departures of a system in the interval $[0, t)$ equals at least the amount data that has left the system in the interval $[0, \tau)$ and the data has been served by the link in $[\tau, t)$. In case τ points at the beginning of a backlog period, i.e. there is no data to be served at time τ, the

cumulative departures equal the cumulative arrivals, i.e. $D(\tau) = A(\tau) = 0$. Thus, it follows that

$$D(t) \geq A(\tau) + r(t - \tau) \tag{5.2}$$

for any $t \geq \tau$. Unfortunately, the beginning of the last backlog period before t is unknown in general and might be difficult to obtain. However, for $t \geq 0$, equation (5.2) holds for at least one backlog period, i.e. for at least one $\tau \in [0, t]$ and we can specify a lower bound for all departures by

$$D(t) \geq \inf_{\tau \in [0,t]} \left\{ A(\tau) + r(t - \tau) \right\}. \tag{5.3}$$

Generalizing the approach discussed above using a generic service function $S(t)$ that specifies a lower bound on the service experienced by the arrivals of a system reveals the notion of service curves in the network calculus. In the following, we apply a similar notation for the service curve S as for the arrivals A and and departures D stated before. According to Cruz et al. [52] the definition of a deterministic service curve is given by

Definition 1 (Deterministic Min-Plus Service Curve) A non-negative, non-decreasing function $S(t)$ is a deterministic service curve for an arrival process $A(t)$ if the corresponding departure process $D(t)$ satisfies for all $t \geq 0$.

$$D(t) \geq \inf_{\tau \in [0,t]} \left\{ A(\tau) + S(t - \tau) \right\} =: A \otimes S(t). \tag{5.4}$$

Here, the operator $\otimes$ denotes the min-plus convolution.

Given the link also delays the departures at most by T, $S(t) = R(t - T)^+$ with $(x)^+ = max\{0, x\}$ describes a latency-rate service curve. Latency-rate service curves can be used to model packetized fair scheduling algorithms and are well known for a number of scheduling disciplines such as weighted fair queuing [53, 78, 120].

Comparing min-plus systems theory to classical systems theory, where the convolution of arriving signal and impulse response holds with equality, equation (5.5) only yields a lower bound for the departure process. However, it is also possible to derive an upper bound, referred to as upper service curve. In case the lower service curve equals the upper service curve and equation (5.5) holds with equality, the resulting function is referred to as exact service curve.

Closely related is max-plus algebra, see, e.g., [21, 43], which translates to min-plus network calculus if the service curve is inverted from a function of packets into a function of time [43]. We denote the arrival and departure timestamps of a packet with index $n \in \mathbb{N}$ by $a(n)$ and $d(n)$, respectively. Hence, the time between the arrival of packet ν and packet n can be described by $a(\nu, n) = a(n) - a(\nu)$. Similar to equation (5.5), we relate the departures of a max-plus linear and shift invariant system to the arrivals as follows

Definition 2 (Deterministic Max-Plus Service Curve) A non-negative, non-decreasing function $s(n)$ is a deterministic service curve for an arrival process $a(n)$ if the corresponding departure process $d(n)$ satisfies for all $t \geq 0$.

$$d(n) \geq \max_{\nu \in [0,n]} \{a(n) + s(n - \nu)\} =: a * s(n) \tag{5.5}$$

Here, $*$ denotes the max-plus convolution.

The deterministic network calculus, however, only allows for a worst-case analysis which yields either overly pessimistic performance bounds, especially if aggregated traffic comes into account, or even no performance bounds for systems, like wireless networks, that do not provide any deterministic service guarantees. To overcome these limitations, recent stochastic network calculus is developed, see, e.g., [47, 58, 80] and references therein. Here, the random service of a system can be modeled by so-called ε-effective

service curves that describe service guarantees which are violated at most with probability ε. Hence, a stochastic max-plus service is given by

Definition 3 (Stochastic max-plus service curve) Consider a system with packet arrival and departure times $a(n)$ and $d(n)$ respectively. The system has a stochastic max-plus service curve $s_\varepsilon(n)$ with violation probability ε if for all $n \geq 1$

$$\mathsf{P}\left[d(n) \geq \max_{v \in [0,n]} \{a(n) + s_\varepsilon(n - v)\}\right] \geq 1 - \varepsilon.$$

Given a stochastic max-plus service curve it is straightforward to compute packet delays defined as $d(n) - a(n)$ that are violated at most with probability ε from $a * s_\varepsilon(n) - a(n)$, e.g., to determine the playout delay of a video application. Furthermore, the stochastic network calculus facilitates analyzing wireless networks. It is used, e.g., in [57, 80, 113, 145] to model physical-layer effects that are due to fading and interference.

5.2 THE DCF VIEWED AS A GPS EMULATION

As mentioned above, GPS is a fluid-flow model that defines a weighted fair resource allocation. Each flow indexed i is assigned a weight φ_i. Considering only backlogged flows, flow i is guaranteed a share of $\varphi_i / \sum_j \varphi_j$ of the capacity C. Due to the granularity of packets real implementations can only emulate GPS with limited precision. Analytical models specify the deviation from an ideal GPS system using worst-case error terms. A prominent model is the Guaranteed Rate Clock (GRC) [61] that contributes the basis of Integrated Services [37]. Here, the error terms can be calculated analytically and are well known for a wide range of schedulers. For an overview we refer to, e.g., [61, 78, 136] and for enhancements of the DCF [23, 76, 97, 141]. By showing that guaranteed rate schedulers with rate r and error term e offer latency-rate service curves $S(t) = R(t - l_{max}/R - e)^+$, [100]

establishes a close link between the GRC model, which is phrased in max-plus algebra, and latency-rate service curves in the deterministic min-plus network calculus This is elaborated further in [78, 138].

In the sequel, we derive a related model for the DCF that we refer to as the *DCF Clock*.

Lemma 1 (DCF Clock) *Consider M stations with indexes i that contend for the medium using the DCF. Let $a_i(n)$ be the arrival time of the n-th packet at station i with length L_i. If the medium is busy at $a_i(n)$ let δ be the residual service time of the packet in service or else δ = 0. The departure times are*

$$d_i(n) = \max\{a_i(n) + \delta, d_i(n-1)\} + \frac{L_i}{r_i} + \phi_i(n) + \psi_i(n)$$

where the average service rate

$$r_i = \frac{L_i}{(1 + E[R_i]) \left(C(\mu + M\Delta) + \sum_{j=1}^{M} L_j \right)} C$$

is subject to zero mean error terms

$$\phi_i(n) = \sum_{m=1}^{1+R_i} b_m(n) - (1 + E[R_i])\mu$$

and

$$\psi_i(n) = \sum_{m=1}^{1+R_i} \sum_{j=1}^{M} K_j(n) \left(\frac{L_j}{C} + \Delta \right) - (1 + E[R_i]) \sum_{j=1}^{M} \left(\frac{L_j}{C} + \Delta \right).$$

Here, C is the capacity, $b_j(n)$ are i.i.d. exponential countdown values with mean μ that are independent of the number of retransmissions R_i, $K_j(n)$ are the inter-transmissions, and Δ comprises all constant per-packet protocol latencies.

Proof   For the first case, we assume that station *i* is not backlogged. If the medium is idle, station *i* starts the medium access procedure immediately at $a_i(n)$ or else, due to carrier sensing, after the residual service time of the packet that is in service, i.e. at $a_i(n) + \delta$. Otherwise, if station *i* is backlogged

it starts the access procedure for packet n after packet $n-1$ finishes service, that is at $d_i(n-1)$. Combining all cases station i initiates its access procedure at $\max\{a_i(n) + \delta, d_i(n-1)\}$.

Before station i transmits packet n an amount of channel idle time of $b_i(n)$ has to be accumulated to complete the countdown procedure. In parallel, all other backlogged stations perform their countdown procedure to contend for the medium.

The transmission of a packet under the DCF includes constant protocol overhead for DIFS, preamble, SIFS, and acknowledgement that are summed up in Δ. Hence, it takes $L/C + \Delta$ units of time to transmit a packet of length L on a channel with capacity C. The number of packets transmitted by station j in the interval $[\max\{a_i(n) + \delta, d_i(n-1)\}, d_i(n)]$ is denoted $K_j(n)$. Thus, the transfer of all packets takes $\sum_{j=1}^{M} K_j(n)(L_j/C + \Delta)$ units of time.

In case of a collision, a retransmission occurs and the transmission procedure has to be applied again. Hence, assuming a random number of retransmissions R_i per transmission attempt, a station has to wait for an additional amount of channel idle time $\sum_{m=1}^{R_i} b_m(n)$ due to the backoff process as well as for an additional time $\sum_{m=1}^{R_i} \sum_{j=1}^{M} K_j(n)\left(\frac{L_j}{C} + \Delta\right)$ due to inter-transmissions.

Assembling all parts the overall departure time follows by

$$d_i(n) = \max\{a_i(n) + \delta, d_i(n-1)\} +$$
$$\sum_{m=1}^{1+R_i}\left(b_m(n) + \sum_{j=1}^{M} K_j(n)\left(\frac{L_j}{C} + \Delta\right)\right). \tag{5.6}$$

In the following, we show that the two error terms have zero mean. To this end, based on the assumption that backoff values and the number of retransmissions are independent, a moment of consideration reveals that $E[\phi_i(n)] = E[\sum_{j=1}^{1+R_i} b_i(n) - (1 + E[R_i])\mu] = 0$[1]. Moreover, we instantiate Theorem 2 with $M = 2$ and $l = 1$ to find the number of inter-transmissions of one station. We have $E[K_j(n)] = 1$ such that $E[\psi_i(n)] = 0$.

---

1 A formal proof can be found in the Appendix on page 107

The residuum after substitution of the error terms in equation (5.6) is the mean latency caused by the countdown procedure and by inter-transmissions. We equate this latency with L_i/r_i where r_i has the interpretation of an average service rate. Hence, we obtain

$$(1 + E[R_i]) \left(\mu + \sum_{j=1}^{M} \left(\frac{L_j}{C} + \Delta \right) \right) = \frac{L_i}{r_i}$$

and solve for r_i to derive the average service rate.   ∎

Lemma 1 specifies the deviation $\phi_i + \psi_i$ of the DCF clock from an ideal GPS system with rate allocation r_i. Note that the error terms have zero mean such that r_i is the true average service rate. The service rate considers the resource consumption that is due to protocol overhead. Apart from that, the rate allocation is proportional to the packet lengths used by individual stations, which formally derives from the target of packet-level fairness of the DCF. Roughly speaking, the packet lengths L_i take the place of the GPS weights φ_i.

It is worthwhile comparing the error terms in Lemma 1 with the GRC model [61] mentioned before. The GR Clock is defined as

$$GRC_i(n) = \max\{a_i(n), GRC_i(n-1)\} + \frac{L_i}{r_i}$$

where departures are subject to an error term χ such that

$$d_i(n) \leq GRC_i(n) + \chi.$$

Compared to the GR Clock the recursion in Lemma 1 uses the actual departure times d_i instead of the target GRC_i. As a consequence, the per-packet error terms of the DCF are accumulated during a busy period. This is not the case in the GRC model. In other words, a GRC scheduler that deviates from the GR Clock, nevertheless has to keep up with the GR Clock at subsequent packet transmissions, i.e. unfairness cannot accumulate. In

contrast the DCF does not seek to correct previous deviations, i.e. the DCF is memoryless in the sense that it does not compensate past unfairness. Moreover, the error terms of known GRC implementations are typically small deterministic upper bounds, e.g., $\chi = L_{\max}/C$ for WFQ [61, 78] as opposed to the DCF error terms that are random and possibly unbounded.

5.3 A STOCHASTIC SERVICE CURVE MODEL OF THE DCF

In this section we derive a stochastic service curve model for the DCF that we phrase in max-plus algebra. Given a stochastic max-plus service curve it is straightforward to compute packet delays defined as $d(n) - a(n)$ that are violated at most with probability ε from $a * s_\varepsilon(n) - a(n)$, e.g., to determine the playout delay of a video application.

Theorem 3 (DCF service curve) *Assume Lemma 1, let all packets have the same size denoted L, and consider a tagged flow. The DCF has a stochastic latency-rate service curve*

$$s_\varepsilon(n) = T + \frac{n}{r}$$

with latency T and rate r defined as

$$T = \tau + (1 + \alpha + \varsigma)\left(\frac{L}{C} + \Delta\right) \quad and \quad \frac{1}{r} = \vartheta + (1 + \beta + \rho)\left(\frac{L}{C} + \Delta\right)$$

and violation probability $\varepsilon = \sum_{m=1}^{\infty}(\varepsilon_1(m) + \varepsilon_2(m) + \varepsilon_3(m))$. Parameters α, β, τ, ϑ, ς, $\rho \geq 0$ and ε_1, ε_2, ε_3 are defined in equations (5.8), (5.9), and (5.10).

The service curve in Theorem 3 is an affine function that comprises a latency offset T and a packet rate r, respectively per-packet latency r^{-1}. The terms correspond to the latency-rate service curve model in min-plus algebra. The free parameters define the service guarantee and determine its violation probability. The parameter choice is subject to numerical optimization. We find that τ, ϑ that stem from the variable duration of the countdown

procedure have comparably small impact, whereas ς, ρ that consider the random amount of time consumed by inter-transmissions as well as α, β that are related to the number of retransmissions have a significant effect.

Proof We consider station $\hat{M}$ and denote its arrivals $a(n)$ and departures $d(n)$. We analyze a single busy period starting at $a(m)$. Applying Theorem 2 to derive the number of inter-transmissions K and from equation (5.6) we obtain by recursion that

$$d(n) = a(m) + \delta + \sum_{i=m}^{n+R} b(i) + (l + R + K)\left(\frac{L}{C} + \Delta\right) \tag{5.7}$$

where $l = n - m + 1$ is the number of packets sent by station $\hat{M}$, $K = \sum_{j=1}^{M-1}\sum_{i=m}^{n+R} K_j(i)$ is the sum of all inter-transmissions since the start of the busy period, and R is the number of retransmissions.

First, we derive probabilistic affine upper envelopes for the random terms. $\sum_{i=m}^{n+R} b(i)$ is the random sum of $l + R$ i.i.d. exponential random variables each with mean μ. Hence, the sum is Gamma distributed and has the well-known moment generating function $M_{\sum b}(\theta, l') = (1/(1 - \theta\mu))^{l'}$ for $\theta < 1/\mu$ with $l' = l + R$. From Chernoff's bound we obtain

$$P\left[\sum_{i=m}^{n+R} b(i) \geq \tau + \vartheta l'\right] \leq e^{-\theta(\tau + \vartheta l')}\left(\frac{1}{1 - \theta\mu}\right)^{l'} \quad \forall \theta \in (0, 1/\mu).$$

Minimization yields $\theta = (\tau + (\vartheta - \mu)l')/((\tau + \vartheta l')\mu)$ such that $P\left[\sum_{i=m}^{n} b(i) \geq \tau + \vartheta l'\right] \leq \varepsilon_1(l')$ where

$$\varepsilon_1(l') = \left(\frac{\vartheta + \frac{\tau}{l'}}{\mu} e^{-\frac{\mu + \vartheta + \frac{\tau}{l'}}{\mu}}\right)^{l'}. \tag{5.8}$$

For the random number of retransmissions R we assume the decoupling approximation introduced by Bianci [31] and calculate a fixed collision probability p_c that remains constant for each transmission attempt and is equal for all stations. Hence, the transmission opportunity can be modeled by a Bernoulli experiment and the probability that a station successfully

transmits a certain number of packets l equals the l-th success in the $(l + R)$-th Bernoulli trial. As for the number of inter-transmissions proven before, this event is negative binomially distributed. We use Corollary 3 to bound the number of inter-transmissions K as well as the number of retransmissions R. We let $r = \alpha + \beta l$ to find $P[R \geq \alpha + \beta l] \leq \varepsilon_2(l)$ where

$$\varepsilon_2(l) = \left(\frac{p_c \, (1 - p_c)^{\beta + \frac{\alpha}{l}} \, (1 + \beta + \frac{\alpha}{l})^{1 + \beta + \frac{\alpha}{l}}}{(\beta + \frac{\alpha}{l})^{\beta + \frac{\alpha}{l}}} \right)^l. \tag{5.9}$$

Similarly, for the inter-transmissions, we let $k = \varsigma + \rho l'$ with $l' = l + R$ to find $P[K \geq \varsigma + \rho(l')] \leq \varepsilon_3(l')$ where

$$\varepsilon_3(l') = \left(\frac{p \, (1 - p)^{\rho + \frac{\varsigma}{l'}} \, (1 + \rho + \frac{\varsigma}{l'})^{1 + \rho + \frac{\varsigma}{l'}}}{(\rho + \frac{\varsigma}{l'})^{\rho + \frac{\varsigma}{l'}}} \right)^{l'}. \tag{5.10}$$

Next, we use Boole's inequality [127] which states that for any finite set of events, the occurrence probability of at least one single event does not exceed the sum of all event risks. To this end, we sum $\varepsilon_i(l)$ over all l to derive a corresponding sample path bound [47]. We estimate the tail probabilities to verify that this sample path bound exists. From equation (5.8) we find

$$\varepsilon_1(l') \leq \left(\frac{\vartheta}{\mu} e^{-\frac{\mu + \vartheta}{\mu}} \right)^{l'} = q_1^{l'}$$

where $q_1 < 1$ generally[2]. Thus, it is evident that $q_1^{l'} \leq q_1^{l}$ and therefore $\varepsilon_1(l') \leq \varepsilon_1(l)$. Thus, we can relax the affine upper envelope and use $\tau + \vartheta l$ instead of $\tau + \vartheta l'$ which simplifies the formulation of an affine upper service curve in the following, however, to the expense of a slightly loser bound. Similarly, we have from equation (5.9) that

$$\varepsilon_2(l) \leq \left(\frac{p_c \, (1 - p_c)^{\beta} \, (1 + \beta)^{1 + \beta}}{\beta^{\beta}} \right)^l = q_2^l$$

---

2 A formal proof for $q_1 < 1$ can be found in the Appendix on page 108.

where $q_2 < 1$ can be shown[3] for $\beta > (1 - p_c)/p_c$. From equation (5.10) we find that

$$\varepsilon_3(l') \leq \left(\frac{p\,(1-p)^\rho\,(1+\rho)^{1+\rho}}{\rho^\rho} \right)^{l'} = q_3^{l'}$$

where $q_3 < 1$ for $\rho > (1-p)/p = M - 1$. For $q_3 < 1$ it is evident that $\varepsilon_3(l') \leq \varepsilon_3(l)$. Thus, we can relax the affine upper bound using $\varsigma + \rho l$ instead of $\varsigma + \rho l'$ as well.

Last, using the geometric sum, we find that

$$\sum_{l=n}^{\infty} \varepsilon_i(l) \leq \frac{q_i^n}{1 - q_i}$$

proving the boundedness of the violation probability. Inserting the envelopes $\alpha + \beta l$, $\tau + \vartheta l$, and $\varsigma + \rho l$ into equation (5.7) and bounding δ by $\Delta + L/C$ yields that

$$d(n) \leq a(m) + \tau + \vartheta l + (1 + \alpha + \varsigma + (1 + \beta + \rho)l)\left(\frac{L}{C} + \Delta \right) \qquad (5.11)$$

is violated at most with probability $\varepsilon_1(l) + \varepsilon_2(l) + \varepsilon_3(l)$.

In the final step we take the maximum over all m of the right-hand side in equation (5.11) to resolve the assumption that the busy period starts at $a(m)$ [43, 101] and obtain the max-plus convolution form in Definition 3. After some reordering equation (5.11) yields the service curve in Theorem 3. ∎

We provide an example of the service curve for IEEE 802.11g. We set $\Delta = 0.1$ ms and $C = 54$ Mbps. We consider $M = 2$ stations and packets of size $L = 1500$ Byte. The effects of the parameters $\tau, \vartheta, \alpha, \beta,$ and ς, ρ on the cumulative error terms are shown in Figure 16. We find that the parameters τ, ϑ have comparably small impact and $\tau = 1.5$ ms and $\vartheta = 0.1$ ms achieve already $\sum \varepsilon_1 < 10^{-6}$. Moreover, we consider $\alpha = 5$ and $\beta = 2$ where we

---

3 A formal proof for $q_3 < 1$, that also holds for $q_2 < 1$, can be found in the Appendix on page 109.

Figure 16: The free service curve parameters τ, ϑ, and α, β, and ς, ρ define sample path bounds for the backoff waiting time (16a, 16b), the number of retransmissions (16c, 16d), and the number of inter-transmissions (16e, 16f) respectively. These bounds are violated at most with probability $\sum \varepsilon_1$, $\sum \varepsilon_2$, and $\sum \varepsilon_3$. The parameters have significant impact on the shape of the resulting latency-rate service curve.
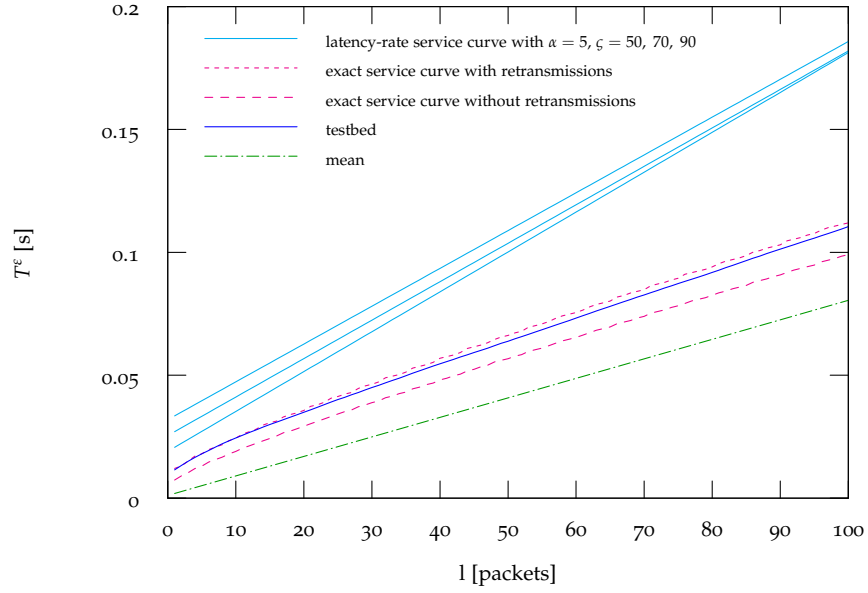
Figure 17: Service curve examples for IEEE 802.11g WLAN illustrating the service experienced by packet trains of length l. We compare ε-effective latency-rate service curves derived from our model to service curves that use numerical solutions of the random terms instead of approximations[4]. In addition we present the mean as well as the 0.99999 service times measured in our testbed. We find that the analytical results match the testbed data quite well.

have $\sum \varepsilon_2 \approx 3 \cdot 10^{-6}$ as well as $\varsigma = 50$ and $\rho = 1.5$ where Figure 16 reveals $\sum \varepsilon_3 \approx 1 \cdot 10^{-6}$. The corresponding stochastic latency-rate service curve is $s_\varepsilon(n) = 19.5 + 1.55n$ ms where $\varepsilon = \sum \varepsilon_1 + \sum \varepsilon_2 + \sum \varepsilon_2 \approx 5 \cdot 10^{-6}$.

Figure 17 depicts the service, i.e. the cumulative delay, experienced by packet trains of length l in an IEEE 802.11g wireless LAN with $M = 2$ stations. It compares various examples of ε-effective latency-rate service curves derived from our model with $\varepsilon = 5 \cdot 10^{-6}$ to service curves that apply numerical solutions of the random terms in equation (5.7) instead of the approximations, i.e. Chernoff's bound and Bool's inequality, used in Theorem 3. Additionally, we present the 0.99999 service times, which state that 0.99999 of all packet trains have been transmitted in less than T^ε seconds, and the mean service rate, which is related to the mean data rate. To this end, we analyzed 10^9 packets that were captured in our testbed from contending greedy stations over a period of approx. 10 days. We find

---

4 In this figure, we call these service curves *exact service curves* since the computation method is exact. However, please note that this is not the standard definition. As stated before, the
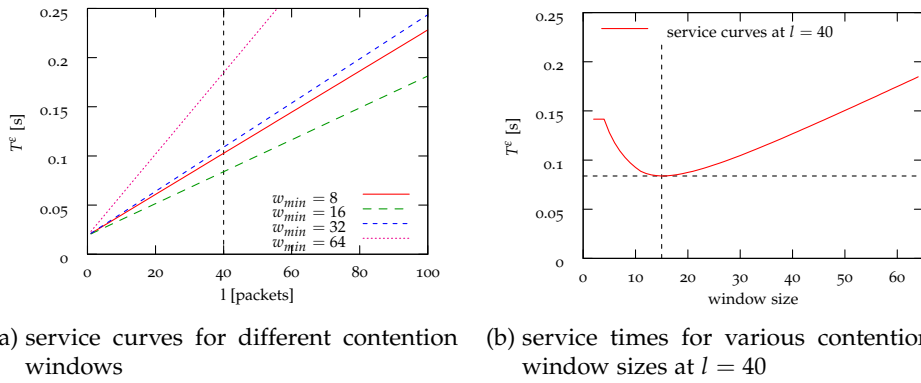
(a) service curves for different contention windows

(b) service times for various contention window sizes at $l = 40$

Figure 18: Service curves with $\alpha = 5$ and $\varsigma = 50$ for various contendion window sizes w. We can derive the optimal contention window size, with respect to the minimal packet burst delay, by finding the lowest service curve.

that the analytical results provide upper bounds to the measured ones and match the testbed data quite well.

Besides the applications within the stochastic network calculus, e.g. to derive end-to-end backlog and delay bounds, we may use the stochastic service curve calculated above to derive the optimal initial contention window size with respect to a minimal packet delay. To this end, we calculate latency-rate service curves for different minimal contention window sizes. Figure 18a illustrates service curves for minimal contention window sizes $w = 8, 16, 32, 64$ exemplarily. Moreover, Figure 18b depicts the service times for packet bursts of $l = 40$ packets over contention window sizes $w = 1, ..., 64$ that are derived from the corresponding service curves. We find that the service curve with $w = 16$ yields the lowest service times for all burst size and all contention windows sizes, respectively. We conclude, that, according to Table 1 in Chapter 2, the IEEE 802.11a and IEEE 802.11g standards already use the best possible inital contention window size.

In the following chapter we view the available bandwidth, more precisely the achievable bandwidth, as a fraction of service experienced by a system. That is, the available bandwidth is directly related to the slope of a latency-rate service curve.

---

network calculus actually refers to exact service curves if the lower service curve equals the upper service curve and equation (5.5) holds with equality.

WIRELESS BANDWIDTH ESTIMATION

In Chapter 6, we focus on a subset of services, namely the available bandwidth. To this end, we report results from an extensive measurement study of active probing methods to shed light on issues such as the specific wireless network characteristics that influence wireless bandwidth estimation, the optimal probing method, and optimal probing parameters. We performed measurements of an IEEE 802.11g network in a highly controlled local wireless testbed that was located in a shielded and anechoic measuring room. Based on the fair queuing model, explored in Chapter 4, we show that iterative methods, which are based on the FCFS assumption, can be expected to estimate the fair share of a new flow instead of the available bandwidth. Our measurements, using the methods in Table 2, support the anticipated results. Furthermore, the gathered data reconfirms the known dependency of bandwidth estimates on the probing packet size.

Equipped with the findings in Chapter 4, we present our own probing method, referred to as *WiProbe*, and show how Kalman filtering can be used to improve bandwidth estimates. We model the fair bandwidth share in wireless networks as a random time-varying process that is at least piecewise stationary. We sample this process continuously to adapt to changes in the channel state, e.g., due to a changing number of contending stations, rate adapting stations, or varying cross-traffic, using probing trains. We show that the estimate variance is related to the train length and present a procedure for tuning the Kalman filter based on analytical findings. Moreover, we elaborate the relationship between smoothness and agility of the Kalman estimates, on the on hand, and the train length and probing interval, on the other. We substantiate our findings by an extensive set of measurements.

Table 2: Bandwidth estimation tools that are considered in this thesis

Probing method		Probing traffic	Inference technique
Pathload	[71]	packet trains	iterative
Pathchirp	[125]	packet chirps	iterative
Spruce	[137]	packet pairs	direct
IGI	[67]	packet trains	direct
PTR	[67]	packet trains	iterative
WBest	[105]	packet trains	direct
DietTOPP	[81]	packet trains	iterative
BART	[54]	packet trains	iterative
SCest	[111]	packet trains	iterative
WiProbe	[36]	packet trains	direct

6.1 METHODS FOR AVAILABLE BANDWIDTH ESTIMATION

In this section we review state-of-the-art measurement-based bandwidth estimation methodologies in wired and wireless networks. We focus on publicly available bandwidth estimation tools, see Table 2, which are also used for measurements in Section 6.3. For related empirical evaluations in wired networks we refer to, e.g., [135, 137].

The task of available bandwidth estimation is to infer the portion of the capacity of a link or a network path that remains unused by cross-traffic. The available bandwidth of a link with index i can be defined as [71]

$$AvBw_i(\tau, t) = C_i(1 - u_i(\tau, t)) \tag{6.1}$$

where C_i is the capacity and $u_i \in [0, 1]$ is the utilization by cross-traffic in the interval $[\tau, t)$. The available bandwidth of a network path is determined by the available bandwidth of the tight link as $AwBw(\tau, t) = \min_i\{AvBw_i(\tau, t)\}$, see, e.g., [71].

Active measurement methods inject specific probes into the network and estimate the available bandwidth from measurements of the probing traffic at the ingress and at the egress of the network. The majority of the methods

uses packet pairs, i.e. two packets sent with a defined spacing in time referred to as gap, or packet trains, i.e. a larger number of packets sent at a defined constant rate. The rate of a packet train can be converted into a certain spacing of the train's packets, showing a direct relation to the gap model of packet pairs. Packet chirps [125] are specific packet trains that are sent at a geometrically increasing rate respectively with a geometrically decreasing gap.

Many methods use a simplified network model, where cross-traffic is viewed as constant rate fluid and the network is abstracted as a single tight link. Under these assumptions the available bandwidth of a network path simplifies to $AvBw = C(1 - u)$. In addition, FCFS multiplexing is usually assumed, where flows share the capacity of a link proportionally to their offered rates. For constant rate probes an expression referred to as rate response curve [108, 114] can be derived as

$$\frac{r_{in}}{r_{out}} = \max\left(1, \frac{r_{in} + \lambda}{C}\right) = \begin{cases} 1 & , \text{if } r_{in} \leq C - \lambda \\ \frac{r_{in} + \lambda}{C} & , \text{if } r_{in} > C - \lambda \end{cases} \tag{6.2}$$

where r_{in} and r_{out} are the input and output rates of probes respectively and λ is the input rate of cross-traffic. If $\lambda \leq C$ the available bandwidth follows as $AvBw = C - \lambda$ and otherwise $AvBw = 0$. Based on this model the task of available bandwidth estimation is to select the rate of probing traffic such that equation (6.2) can be solved for C and λ or $C - \lambda$. While equation (6.2) is usually used for packet train probes, an equivalent gap response curve can be derived for packet pairs, where the gap g is linked to the rate r by the packet size l resulting in $g_{in} = l/r_{in}$ and $g_{out} = l/r_{out}$ [108].

In [72] measurement methods are classified by their inference technique as either *direct* or *iterative* probing schemes. Direct probing schemes often assume that the capacity of the link C is known in advance. In this case equation (6.2) can be solved for the rate of the cross-traffic if the probing rate is larger than the available bandwidth. A straightforward choice is to

probe with $r_{in} = C$ in which case the available bandwidth follows from equation (6.2) in rate respectively gap notion as [105, 137]

$$AvBw = C \left(2 - \frac{C}{r_{out}} \right) = C \left(1 - \frac{g_{out} - g_{in}}{g_{in}} \right). \qquad (6.3)$$

Spruce, WBest, and IGI are methods that use direct probing. Spruce assumes that the capacity is known a priori and immediately applies the gap version of equation (6.3). WBest provides a two-step algorithm using packet pairs to estimate the link capacity and packet trains for available bandwidth estimation based on the rate version of equation (6.3). IGI uses probing trains with increasing gaps resulting in a more complex direct probing formula than equation (6.3). For more details see [67].

Iterative probing methods do not require a-priori knowledge of the link capacity. They employ an iterative procedure with multiple probing rates aiming to locate the turning point of the rate response curve equation (6.2), i.e. they seek to find the largest probing rate r_{in} such that $r_{in}/r_{out} = 1$. At this point the probing rate coincides with the available bandwidth.

TOPP, DietTOPP, BART[1], PTR, SCest, Pathload, and Pathchirp are iterative probing methods. TOPP [114] uses trains of packet pairs with increasing rate and applies equation (6.2) for available bandwidth estimation. It recursively extends the model to the multiple node case and in addition it estimates the capacity from the second linear segment in equation (6.2). Closely related is a simplified version called dietTOPP. The BART tool, a successor of DietTOPP uses a Kalman filter to obtain both, the end-to-end available bandwidth and the bottleneck capacity. It infers the turning point of the rate response curve using a linear system that describes the packet dispersion by two parameters of the sloping straight line in the overload region. To this end, BART uses packet trains with a probing rate greater than the available bandwidth, i.e., $r_{in} > AvBw$ and applies a two-dimensional Kalman filter to estimate these unknown parameters. The resulting system of equations,

---

1 BART is not publicly available. Hence, it is omitted by our measurements in Section 6.3.

however, is quite complex and the authors do not provide a way for tuning the filter parameters optimally. PTR is a packet train method that uses a gap version of equation (6.2). Pathload varies the rate of packet trains using a binary search algorithm to find the largest probing rate that does not cause overload and hence matches the available bandwidth. It uses increasing one-way delays as an indication of overload. Increasing delays indicate that the input rate exceeds the output rate, i.e. $r_{in}/r_{out} > 1$ which clearly shows the relation to equation (6.2) [108]. Pathchirp increases the probing rate within a single packet train, referred to as a chirp, instead of varying the rate of successive packet trains. Like Pathload it detects the turning point of the rate response curve from increasing one-way delays. SCest makes use of novel measurement methodology based on network calculus as well as findings in [107] and interprets bandwidth estimation as the inversion problem of $D = A * S$ for S of max-plus linear systems. Hence, it tries to infer not only the available bandwidth but the complete service, represented by a service curve, which is offered by a system. Transferring the problem to max-plus algebra, the probing tool operates on packet time stamps and computes a stochastic ϵ-effective max-plus service curve from steady-state delay percentiles obtained from probing packets.

Most of the discussed methods have been developed for wired networks, while WBest and dietTOPP have been suggested by the authors for available bandwidth estimation in wireless networks. SCest is stated to be used in lossy systems with arbitrary scheduling disciplines and hence, should be suitable for wireless systems as well. Further on, a method called Probe-Gap has been proposed for bandwidth estimation in broadband access networks [98]. The method does not exactly fit into the classification scheme used here. ProbeGap sends out single packets and collects the one-way delays of these probes. The fraction of the packets which have a delay close to zero are assumed to have found an idle channel. This fraction is used to estimate the available bandwidth. Besides, passive measurement approaches can take advantage of the wireless broadcast medium [130]

or protocol-related information [102]. Passive methods are, however, not considered within the scope of this thesis.

6.2 AVAILABLE BANDWIDTH ESTIMATION IN WIRELESS SYSTEMS

In the following, we discuss the influence of relevant wireless link characteristics, discussed in Chapter 4, and the distributed medium access procedure, that is described in Chapter 2, to state-of-the-art bandwidth estimation techniques. We show how these aspects affect current fluid rate and gap models and reason which quantity we expect to be estimated by known methods for bandwidth estimation in wireless systems.

Owing to our findings in Chapter 4, the long-term fair share f at a congested wireless link can be computed as the solution of

$$f : \sum_{k=1}^{n} \min\{r_k, f\} = C \tag{6.4}$$

where C is the capacity[2], r_k is the rate of flow k, and n is the number of flows. Once f is determined the output rate of flow k follows as $\min\{r_k, f\}$.

The rate response curve of a fair queuing system follows immediately as

$$\frac{r_{in}}{r_{out}} = \max\left(1, \frac{r_{in}}{f}\right) = \begin{cases} 1 & \text{, if } r_{in} \leq f \\[2mm] \frac{r_{in}}{f} & \text{, if } r_{in} > f \end{cases} \tag{6.5}$$

where r_{in} and r_{out} are the input and output rates of the probes respectively. As opposed to the FCFS rate response curve equation (6.2) the available bandwidth cannot be derived from equation (6.5). Trivially the available bandwidth $AvBw$ is upper bounded by the fair share, i.e. $0 \leq AvBw \leq f$.

---

2 In wireless systems, effects that are due to fading and interference result in a time-varying channel capacity $C(t)$. Thus, it is straightforward to adapt the definition of available bandwidth, see equation (6.1), accordingly. Our controlled measurement environment, however, eliminates these effects to a large extent. Hence, we can assume a constant capacity. This is similar to the fluid rate and gap models for bandwidth estimation, see equation (6.2) and equation (6.3), respectively.

Without further assumptions the two extremal values can, however, be easily attained if $f \leq C/2$. As an example consider a single contending flow with rate $\lambda = C/2$ respectively $\lambda = C$. The fair share of a new greedy flow is $f = C/2$ in both cases, whereas the available bandwidth becomes $AvBw = f$ respectively $AvBw = 0$.

Referring to the classification of bandwidth estimation methods in Section 6.1 we conclude that iterative methods, which use the turning point of the rate response curve as bandwidth estimate, can be expected to report the fair share of a new greedy flow in case of a fair wireless link. For existing direct probing methods that inject probes with rate $r_{in} = C$ such a clear result cannot be established. Inserting $r_{out} = f$ into equation (6.3) does neither compute the available bandwidth nor the fair share. We note, however, that direct probing with $r_{in} = C$ could easily report the fair share, since $r_{out} = f$ in this case.

6.3 EXPERIMENTAL EVALUATION OF BANDWIDTH ESTIMATION

Equipped with the results from Section 6.2 we investigated the performance of the bandwidth estimation tools listed in Table 2. If not mentioned otherwise, we used the default configuration of the bandwidth estimation tools to perform the experiments. We evaluate the methods using a wireless testbed[3] in a shielded, anechoic room. Hence, we assume that the physical medium is free of interference from external sources that do not belong to the testbed. We focus on the accuracy of wireless bandwidth estimates and show how these relate to the available bandwidth respectively to the fair share under different types of contending traffic. We do not report probing

---

3 Basically we used the same testbed configuration with 2 senders and 1 receiver, as already presented in Chapter 4. The hardware, however, was slightly different. We used Lenovo Thinkpad R61 notebooks with 1.6 GHz and 2 GB RAM running Ubuntu Linux 7.10 with kernel version 2.6.22. We employed the internal Intel Pro/Wireless 4965 AG IEEE 802.11 wireless lan adapters. The access point was a Buffalo Wireless-G 125 series running DD-WRT [2] version 24.

overhead, intrusiveness, as well as run or convergence times. These aspects are elaborated, e.g., in [135].

Impact of the Intensity of Contending Traffic

In the first set of experiments we estimate the end-to-end available bandwidth from a sender to a receiver in the presence of a single contending flow. We increase the rate λ of the contending traffic that flows from 0 Mbps up to 28 Mbps in steps of 1 Mbps. The contending traffic consists of packets of 1500 Bytes and is generated using the D-ITG traffic generator [1]. All probe packets are set to 1500 Bytes.

Figure 19 shows the average of 25 available bandwidth estimates for each of the tools and all rates of the contending traffic as well as corresponding confidence intervals at a confidence level of 0.95. As a reference the available bandwidth $AvBw = C - \lambda$ as well as a the fair share of a new flow $f = \max\{C - \lambda, C/2\}$ are plotted, where we use $C = 28$ Mbps from Figure 4 for a packet size of 1500 Bytes.

From our arguments in Section 6.2 we expect that the iterative probing methods Pathload, DietTOPP, Pathchirp, and PTR report an estimate of the fair share. As indicated in Figure 19 the fair share and the available bandwidth are identical for contending traffic with rate $\lambda \in [0 \dots 14]$ Mbps, whereas they differ for $\lambda \in (14 \dots 28]$ Mbps. Figure 19a shows that the estimates from Pathload (which reports an upper and a lower bound of the available bandwidth) and DietTOPP clearly confirm the fair queuing model in equation (6.5). Both methods closely track the fair share and the reported estimates deviate noticeably from the available bandwidth as the rate of the contending traffic increases beyond 14 Mbps. The results from PTR in Figure 19c and to a lesser extent from Pathchirp in Figure 19b confirm this view. In case of Pathchirp, we used the estimates provided after Pathchirp's self-adapting phase. In our experimental results plotted
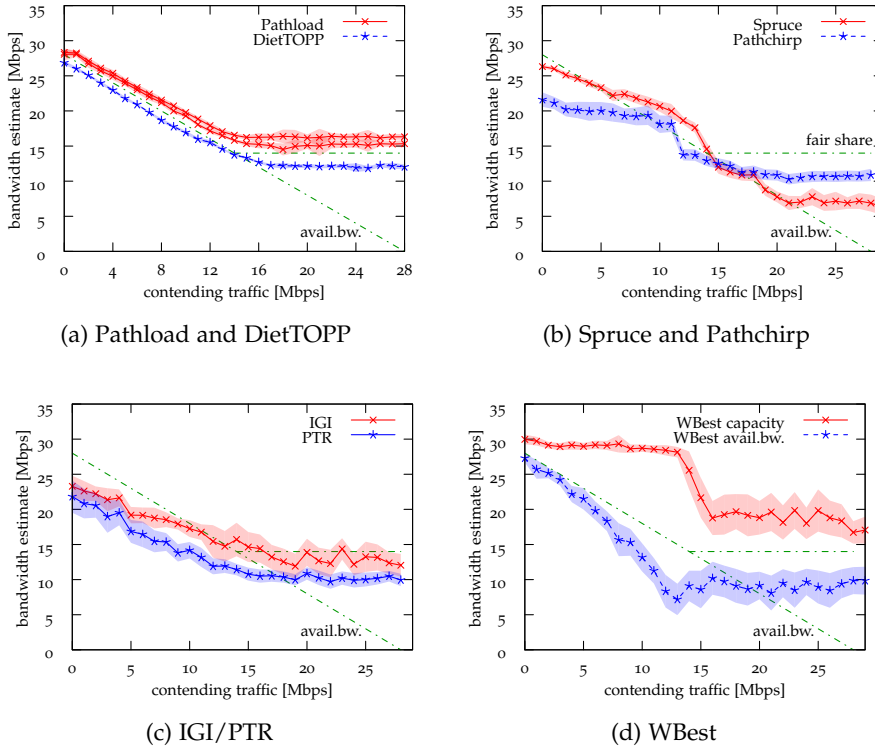
Figure 19: Bandwidth estimates for a wireless link with one contending flow.

in Figure 19b it, however, underestimates the fair share and the estimates exhibit a comparably high variance.

Pathload has been reported to provide inaccurate bandwidth estimates for wireless networks in [34, 98, 105]. This stands in contrast to experiences made by using Pathload in wired networks. In [105] the probing packet size is mentioned as a possible reason for bandwidth underestimation, and [98] identifies the signature of one-way delays in wireless networks as a source of the problem. Having confirmed the strong impact of the packet size on the throughput in wireless networks, see Figure 4, we modified Pathload so that we can specify the probing packet size. Using Pathload with a fixed packet size of 1500 Bytes improves bandwidth estimates significantly. As depicted in Figure 19a, this yields quite accurate and stable results. Similar problems have been reported for IGI/PTR [105], which can also be mitigated using a packet size of 1500 Bytes rather than the default size. The estimates in Figure 19c are, however, less sensitive to the intensity of contending traffic, as also reported for cross-traffic in wired networks in [137].

Direct probing tools require a-priori knowledge of the link capacity. We executed Spruce with a given capacity of $C = 28$ Mbps, which corresponds to the throughput for packets of 1500 Bytes size as shown before in Figure 4. From our results shown in Figure 19b we cannot detect a clear trend of the estimates towards either the fair share or the available bandwidth once λ exceeds 14 Mbps. WBest uses a two-step algorithm to estimate first the capacity and then the available bandwidth. In our measurements both estimates exhibit a comparably high variance as shown in Figure 19d. Moreover, the capacity estimates are sensitive to contending traffic, possibly a result of fair resource allocation, such that bandwidth estimates that are based hereon may be unreliable.

Impact of the Number of Contending Flows

In the second set of experiments we investigate the impact of the number of contending flows on bandwidth estimates. We use contending traffic with a total rate of 20 Mbps, which is divided evenly among 1 to 4 flows. Hence, the available bandwidth $AvBw \approx 8$ Mbps remained constant in all experiments, whereas the fair share of a new flow is approx. $(14, 9.3, 8, 8)$ Mbps for $(1, \ldots, 4)$ contending flows each offering a rate of approx. $(20, 10, 6.6, 5)$ Mbps respectively. Since we only related the estimates of iterative probing methods to the fair share, we restrict the results shown here to iterative methods. Again all contending and probing packets are adjusted to have a fixed size of 1500 Bytes in order to achieve comparability.

As in the experiments presented above, we repeated each experiment 25 times. In Figure 20 we illustrate the average of the available bandwidth estimates and corresponding confidence intervals at a confidence level of 0.95. The estimates of the iterative methods Pathload, DietTOPP, PTR, and Pathchirp are closely related to the fair share and do not match the available bandwidth.

(a) Pathload and DietTOPP   (b) PTR and Pathchirp

Figure 20: Fair share estimates for a wireless link with several contending flows.

The direct probing tools that were investigated as well tend to be more inaccurate and do not report the available bandwidth nor the fair share in these experiments.

6.4 ONLINE ESTIMATION OF FAIR SHARE USING KALMAN FILTERING

In this section we present our own measurement-based rate estimation technique for wireless LANs using the DCF that incorporates the findings regarding the fairness that are presented in Chapter 4. We show how user applications, such as rate-adaptive video streaming, can estimate their fair bandwidth share under the DCF from measurements of their data arrivals and departures. Obviously, this is closely related to the active probing techniques presented above. In contrast to most of the existing methods and tools, we do not assume constant channel and stationary cross-traffic conditions. In oder to detect changes in the channel capacity, we probe continuously over time which requires an accurate and fast probing approach with low probing overhead.

For now we assume an ideal GPS system and a probing flow that transmits a burst of $l + 1$ packets. These packets are marked with time stamps $a(n)$ and $d(n)$ at the sender and the receiver, respectively. To avoid the necessity of synchronized clocks at sender and receiver, we only use the time differences between the packet departures that can be computed from

the time stamps at the receiver only. The average gap of the departures follows from

$$g_d = \frac{d(l+1) - d(1)}{l}$$

and a sample of the fair share follows as $f = L/g_d$.

In a constant fluid system every sample yields an unbiased bandwidth estimate. In case of real scheduling implementations, such as the DCF, these (fair) rate samples can, however, be largely perturbed by the random channel access procedure. Thus, in oder to improve results, most bandwidth estimation tools perform some kind of filtering. For instance, the use of packet trains itself constitute a low-pass filter and most of the tools described above perform some averaging over several samples to obtain more accurate bandwidth estimates. A naive approach for online bandwidth estimation is to employ a moving average over a specific number of past samples. The open question, however, remains a suitable filter parameterization that includes, e.g., train length, probing rate intensity, probing interval, and filter length.

In the following, we view the variability of bandwidth samples as *measurement noise* and denote the changes in the channel capacity as *process noise*. According to our findings in Chapter 4, we model the perturbations within packet train samples as a Gaussian noise process. To this end, we employ a Kalman filter to remove noise from our measurements and perform an online smoothing over the train departure gaps g_d. In case of i.i.d Gaussian process and measurement noises the Kalman filter generates optimal estimates in the sense that it minimizes the mean squared error by recursively weighting past measurements. Even if the Gaussian assumption is dropped, the Kalman filter is still the best linear unbiased estimator. Furthermore, stationarity of the system is not a necessary precondition. This makes the Kalman filter an ideal candidate for estimating the true available bandwidth.

Figure 21: A Kalman filter solves the inversion problem of a linear system. It uses samples (z) of a process observed over time that are perturbed by random measurement noise (v) and other deterministic inaccuracies (u) and calculates new estimates for system states (x) that have a minimal mean squared error.

Generally, we want to estimate the system state vector $x(n)$ at a discrete time n of any linear system, as illustrated in Figure 21, that can be described by a state equation and an output equation respectively

$$x(n) = Ax(n-1) + B(n)u(n) + \omega(n)$$

$$z(n) = Hx(n) + v(n).$$

Here, A is a state transition matrix that describes the dynamic transitions of two time-consecutive states. Additionally, two external inputs are modeled by the framework. The deterministic, i.e. completely known, external input vector $u(n)$ is associated by a control-input matrix B. All random components are viewed as process noise, and represented by a stochastic variable $\omega(n)$ that is assumed to be normally distributed with mean 0 and variance Q, i.e. $u(n) \sim N(0, Q)$. Likewise, the measured output vector $z(n)$, which is linked to real system states through an observation matrix H, is perturbed by a normally distributed measurement noise $v(n) \sim N(0, R)$ as well.

The true system state $x(n)$ cannot be measured directly but obtained by the system observations $z(n)$, which, however, cannot be taken at face values due to the perturbation by noise. To this end, we need to solve the inversion problem of the linear system given above and derive optimal

state estimates by iteratively applying a set of equations, i.e. (6.6) – (6.8), known as the time and measurements updates, as new samples $z(n)$ become available. That is, the Kalman filter and for a general introduction we refer to [45] and references therein.

In our specific use-case, the state estimates, i.e. bandwidth estimates, are scalar and we do not consider any deterministic external inputs. Thus, the transition matrices become scalar as well and it follows that $A = 1$, $B = 0$, and $H = 1$, which simplifies the Kalman equations. To generate smoothed bandwidth estimates of the departure gap we index consecutive measurement samples $z(n) := g_d(n)$ and denote the output of the Kalman filter $x(n) := \hat{g}_d(n)$. Using the Kalman equations [45] we calculate

$$\hat{g}_d(n) = (1 - G(n))\, \hat{g}_d(n-1) + G(n)\, g_d(n) \tag{6.6}$$

where the factor $G(n)$, that is referred to as Kalman gain, represents the impact of the current measurement sample. It is given by

$$G(n) = \frac{\sigma_E^2(n-1) + \sigma_P^2}{\sigma_E^2(n-1) + \sigma_P^2 + \sigma_{g_d}^2} \tag{6.7}$$

where $\sigma_E^2(n)$ is the estimation error variance that acts as a measure for the reliability of the current state estimate $\hat{g}_d(n)$. The bigger σ_E^2 the less reliable the estimate. Typically, the filter is initialized with a guess of the first state estimate $\hat{g}_d(0)$ and a large value $\sigma_E^2(0)$ representing the uncertainty associated with this guess. Updates of σ_E^2 can then be calculated recursively by

$$\sigma_E^2(n) = (1 - G(n))\, (\sigma_E^2(n-1) + \sigma_P^2). \tag{6.8}$$

The two external parameters σ_P^2 and $\sigma_{g_d}^2$ denote the Gaussian process and measurements noise variances, respectively. If these variances are known, the resulting filter is known to be optimal, i.e. no other filter can achieve a smaller mean squared error. For non-Gaussian noises, the filter is still the

best linear estimator. In the following we outline procedures for determining suitable process and measurement noise for bandwidth estimation in case of IEEE 802.11 WLANs using the DCF.

6.5 MEASUREMENT NOISE PARAMETERIZATION

To parameterize the measurement noise, we estimate the variability of the average departure gap g_d. From equation (5.7) we derive that

$$g_d = d(l+1) - d(1) = (l + R + K)\left(\frac{L}{C} + \Delta\right) + \sum_{j=2}^{l+R+1} b(j) \tag{6.9}$$

where K are the inter-transmissions between packet 1 and $l+1$ and R are the corresponding retransmissions. Equation (6.9) contains three sources of randomness: the number of inter-transmissions from contending stations K, the retransmissions R, and the cumulated countdown values $\sum_{j=2}^{l+R+1} b(j)$. As discussed in Chapter 4 the effects due to inter-transmissions and retransmissions are dominant so that we approximate $\sum_{j=2}^{l+R+1} b(j) \approx (l + R)\mu$ by its mean value.

Regarding the inter-transmissions we assume $l \gg 1$ and use the Gaussian approximation[4] from Corollary 2. After normalization we find that K/l is normal with standard deviation $\sigma_{K/l} = \sqrt{(1-p)/(p^2 l)}$ where $p = 1/M$ is determined by the number of contending stations. Similarly, the retransmissions can be approximated by a Gaussian distribution as well and R/l becomes normal with standard deviation $\sigma_{R/l} = \sqrt{(p_c)/((1-p_c)^2 l)}$ where p_c is the collision probability discussed in Chapter 2. Using that $\sigma_{K+R}^2 = \sigma_K^2 + \sigma_R^2$ generally and by dividing equation (6.9) by l we recover g_d and derive the standard deviation of the train gaps as

$$\sigma_{g_d} = \sqrt{\frac{p^2 p_c + (1-p)(1-p_c)}{(1-p_c)^2 p^2 l}} \; (\Delta + L/C). \tag{6.10}$$

---

4 We recall that, according to Section 4.4, we can apply the Gaussian approximation already for packet train length $l > 4$.

For example consider a wireless network with two greedy stations, i.e. $M = 2$. To this end, we find the collision probability $p_c \approx 0.105$ and the channel access probability $p = 1/M = 1/2$. For the constant per packet overhead we assume $\Delta + L/C \approx 32$ ms. Hence, the standard deviation of the measurement noise comes to $\sigma_{g_d} \approx 0.492/\sqrt{l}$ ms.

6.6   PROCESS NOISE PARAMETERIZATION

In addition to the measurement noise, the Kalman filter must also be supplied with the variance σ_P^2 of the measured process. This parameter determines how quickly the filter considers new measurements to be reliable and has implications for the bandwidth estimation process: for example, it determines the convergence speed, influences the tracking capability, and the variability of resulting estimates. However, the variance of the measured process is typically not known in advance. Hence, we choose a value based on the knowledge of the type of changes we are interested in capturing. To this end, we view the adjustment of the process noise as a tuning knob that trades smoothness for convergence speed. We again highlight, that even if the measured process is not Gaussian distributed, the Kalman filter provides the lowest mean squared error achievable by a linear filter.

In order to correctly parameterize our probing technique, it is vital to consider its use-case. In case a user is only interested in the average available bandwidth over several minutes, different settings are required than if changes must be detected within seconds. To this end, we specify the maximum bandwidth change B that can be identified within a given time T_s, denoted as discontinuity period, see Figure 22a. Furthermore, we assume a maximum probing rate that is constrained to a specific average value r_p, resulting in an average inter-train sending time $t_\Delta = L(l+1)/r_p$.

We identify the discontinuity in the bandwidth process by calculating the variance σ_P^2 for a given period T_s after the fair share change. Since the filter works on average train gaps rather than rates, we describe the bandwidth

(a) Influence of process noise on change tracking

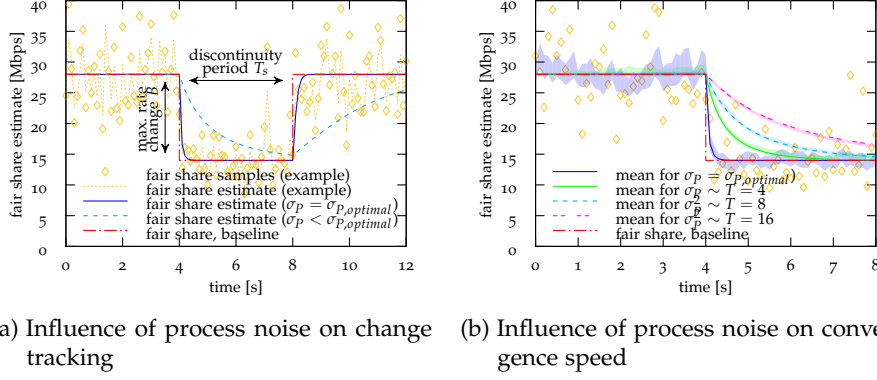(b) Influence of process noise on convergence speed

Figure 22: Averaged simulation results of smoothed Kalman estimates from samples that are corrupted by Gaussian noise. The results clearly show the influence of process noise on change tracking and convergence speed respectively. We find that the smaller the process noise, the smoother the Kalman estimates, however, to the expense of tracking capability and convergence time. For an optimal process noise the filter generates estimates with a minimal mean squared error for a given scenario.

change $B = 14$ Mbps by $g_B = L/B$. As $T_s/t_\Delta = n_p$ samples are generated during T_s, σ_P^2 is calculated as follows[5]

$$\sigma_P^2 = \frac{g_B^2}{n_p} = \frac{g_B^2 L}{T_s r_p}(l+1). \tag{6.11}$$

Evidently, σ_P^2 decreases linearly as the number of samples in the segment is increased. Using equation (6.11) the Kalman filter yields the optimal estimate with respect to a minimal mean squared error within the considered period. The optimality for detecting shorter discontinuities or bigger bandwidth changes respectively can not be guaranteed. As a result, this parametrization can be viewed as a lower bound for the filter tracking ability.

For example, consider the simulation scenario depicted in Figure 22a. After 4 s the fair share abruptly drops from 28 to 14 Mbps which can be considered as a worst-case drop in a 802.11g wireless network when a second greedy stations joins the network. The corresponding gap length for $B = 14$ Mbps is given by $g_B \approx 0.86$ ms assuming a packet length of $L = 1500$ byte. The fair share remains in this state for $T_s = 4$ s. Rate

---

5 The formal derivation of σ_P^2 can be found in the Appendix on page 111.

samples, which are assumed to be collected every $t_\Delta = 100$ ms by probing trains of length $l = 8$, are corrupted by a Gaussian noise process with variance $\sigma_{g_d}^2 \approx 0.03\, ms^2$ according to equation (6.10). To ensure that the filter follows cross-traffic variations, we must set $\sigma_P^2 > 0$. If σ_P^2 is too large, the estimates become unnecessarily noisy. If, on the other hand, σ_P^2 is too small, the discontinuity period is over-smoothed. Thus, we apply equation (6.11) and derive the optimal process noise variance by $\sigma_P^2 \approx 7.4 \times 10^{-8}$. This minimizes the mean squared error for the given scenario and the filter follows the bandwidth change perfectly.

Convergence Speed

We calculate the time the filter needs to converge to a new value after an abrupt bandwidth change. To this end, we examine the step response of the so-called scalar steady-state Kalman filter. From the Kalman equations (6.6) – (6.8) it is evident, that the error variance $\sigma_E^2(n)$ and the Kalman gain $G(n)$ are independent of the current state estimate and the corresponding measurement respectively. Moreover, for stationary noises the parameters $\sigma_E^2(n)$ and $G(n)$ converge to constant values $\lim_{n \to \infty} \sigma_E^2(n) = \sigma_E^2$ and $\lim_{n \to \infty} G(n) = G$. These steady-state values can be calculated analytically offline

$$
\sigma_E^2 = \frac{\sigma_P^2}{2} + \frac{\sigma_P^2}{2}\sqrt{1 + 4\sigma_{g_d}^2/\sigma_P^2}
$$
$$
G = \frac{\sigma_E^2}{\sigma_E^2 + \sigma_{g_d}^2}
$$

which, in general, reduces the computational complexity of the Kalman filter, especially in case of matrix operations.

Next, we substitute the steady-state Kalman gain G in equation (6.6) and find that the steady-state Kalman filter equations are equivalent to the

recursive formulation of the exponentially weighted moving average filter, see, e.g., [18, 45], with smoothing factor G

$$\hat{g}_d(n) = (1 - G)\hat{g}_d(n - 1) + G g_d(n). \tag{6.12}$$

Applying the closed form solution of the exponentially weighted moving average filter, we calculate the step response, that describes an abrupt bandwidth change, by

$$\hat{g}_d = (1 - G)^n u(n) = e^{-\alpha n} u(n).$$

where $u(n)$ is the step function. By substituting $(1 - G) = e^{\alpha}$, we find that the bandwidth converges to a new fair share value exponentially. Using $\cosh(\alpha) = (e^{\alpha} + e^{-\alpha})/2$, the relationship between the noise process variances σ_P^2 and $\sigma_{g_d}^2$ to α, and consequently, G is given by

$$
\begin{aligned}
e^{-\alpha} = (1 - G) \qquad &= 1 + \frac{\sigma_P^2}{2\sigma_{g_d}^2} - \frac{\sigma_P^2}{2\sigma_{g_d}^2}\sqrt{1 + 4\sigma_{g_d}^2/\sigma_P^2} \\
e^{\alpha} = \frac{1}{(1 - G)} \qquad &= 1 + \frac{\sigma_P^2}{2\sigma_{g_d}^2} + \frac{\sigma_P^2}{2\sigma_{g_d}^2}\sqrt{1 + 4\sigma_{g_d}^2/\sigma_P^2} \\
\cosh(\alpha) = 1 + \frac{\sigma_P^2}{2\sigma_{g_d}^2}. &
\end{aligned}
$$

Finally, to calculate the convergence time, we make use of the fact that $e^{-\alpha n}$ generally decays to less than 0.01 of its initial value after $n = 5/\alpha$ steps. Moreover, taking into account that packet train probes are sent every t_Δ seconds on average, we calculate the filter convergence time T for a set of values σ_P^2, $\sigma_{g_d}^2$ by

$$T = 5 t_\Delta / \operatorname{arcosh}(1 + \frac{\sigma_P^2}{2\sigma_{g_d}^2}). \tag{6.13}$$

Estimate Variance

In order to evaluate the influence of the process noise parameterization on the variation of filtered bandwidth estimates, we derive the estimate variance of the filtered estimates during the discontinuity period T_s. To this end we perceive the exponentially weighted moving average filter as an auto-regressive $AR(1)$ process [119] that, has the form $x(n) = c + a_1 x(n - 1) + \varepsilon_t(n)$ and a well known variance $\sigma_{\varepsilon_t}^2 / (1 - a_1^2)$ where $\sigma_{\varepsilon_t}^2$ denotes the variance of a zero-mean white noise process ε_t. By comparing this general form of $AR(1)$ processes to equation (6.12) we find that $c = 0$, $a_1 = 1 - G$, and $\varepsilon_t = G\,x(n)$. Moreover, the variance of the sample process is given by the sum of the process and measurement variance, i.e. $Var(x(n)) = \sigma_P^2 + \sigma_{gd}^2$, which yields $\sigma_{\varepsilon_t}^2 = G^2(\sigma_P^2 + \sigma_{gd}^2)$ for the white-noise process variance. Thus, the overall variance of the filtered estimate process is given by

$$\text{Var}\left(\hat{x}(n)\right) = \frac{G^2(\sigma_P^2 + \sigma_{gd}^2)}{1 - (1 - G)^2}. \tag{6.14}$$

Figure 23 illustrates the relationship between estimate variance and train length for discontinuity periods of different lengths. In case the variance of the cross-traffic is known in advance, we can readily employ equation (6.14) to determine the probing train length, which yields a minimal mean squared error. This is elaborated further in the next section.

Considerations on Probing Train Length

As shown by equation (6.10), increasing the packet train length l reduces the variance of the bandwidth samples proportionally to $1/l$. Additionally, probing bandwidth is wasted when using short packet trains because we use $l + 1$ packets per train. Moreover, the ratio of bandwidth used per train and number of train gaps $(l + 1)L/l$ becomes non-linear. On the other hand, long packet trains, which result in a low sampling rate, have a negative

Figure 23: Estimate variances over probing train length for a 14 Mbit bandwidth change within $T_s = 0.5, 1, 4, ..., 32$ s. The constant average probing rate is approx. 1 Mbps. We find that the estimate variance depends on the measurement noise and therefore packet train length. Thus, we derive the optimal packet train length from the minimal variance for a given scenario.

impact on the bandwidth change tracking accuracy, which is shown by equation (6.11). Hence, we are interested in the optimal train length for a given probing rate r_p, a minimal discontinuity period T_s, and a maximum bandwidth change B.

To obtain the optimal train length, that provides the smallest estimate variance under the given constraints, we need to minimize equation (6.14) in order to minimize the estimate variance. Thus, we find the optimal train length with respect to the parameters listed above as follows[6]

$$l = \frac{\sqrt{L\, T_s\, r_p\, (p^2 p_c^2 + p p_c - p - p_c + 1)}}{L\,(1 - p_c)\, p\, g_B} \,(\Delta + L/C). \qquad (6.15)$$

The train length l is proportional to the square roots of the discontinuity time, the probing rate intensity, and the reciprocate of the bandwidth change, i.e. $l \sim \sqrt{T_s\, r_p / g_B^2}$. This yields some important conclusions: first, equation (6.15) resembles findings, e.g., in [108], which state that infinite long packet trains yield perfect bandwidth estimates in case of constant channel and cross-traffic conditions. For changing conditions, however, using long packet trains yields sub-optimal results.

Another important aspect is the improvement of high probing rates over low probing that is proportional to $\sqrt{r_p}$. Thus, the initial estimation im-

---

6 The formal derivation of the optimal train length can be found in the Appendix on page 112.

provement for low probing rates is significant but becomes less pronounced if the probing rate is increased while the other parameters remain constant. Likewise, the improvement of long over short discontinuity periods decreases.

Figure 23 depicts the variance of the filtered samples for a process noise optimized to track changes of $B = 14$ Mbps within $T_s = [0.5, 1, 2, 4, \ldots, 32]$ s. This variance is related to the mean squared error between the estimate and the actual bandwidth during the discontinuity period of duration T_s. It is comprised of the deviation during the convergence period T and the measurement variance after convergence. It is evident that for the depicted scenario, it is desirable to use short train lengths even if the resulting sample variance is increased: the overall segment mean squared error is minimized. Increasing the probing traffic intensity r_p, however, would yield longer trains that tend to a lower probe variance. This results in better, i.e. smoother, estimates.

6.7 EXPERIMENTAL EVALUATION OF CONTINUOUS BANDWIDTH ESTIMATION

To demonstrate our filtering approach, we implemented a modular, portable measurement framework, called WiProbe. It is based on the findings in the previous sections and implements direct probing for estimating the fair share in wireless networks. Kalman filtering is used to continuously remove measurement noise from the probes. In order to evaluate our method, we performed experiments in a controlled testbed environment containing both – wired and wireless links. We investigate the performance of the Kalman filter and the effects of parametrization to provide an underpinning to our theoretical findings in Section 6.5 and Section 6.6.
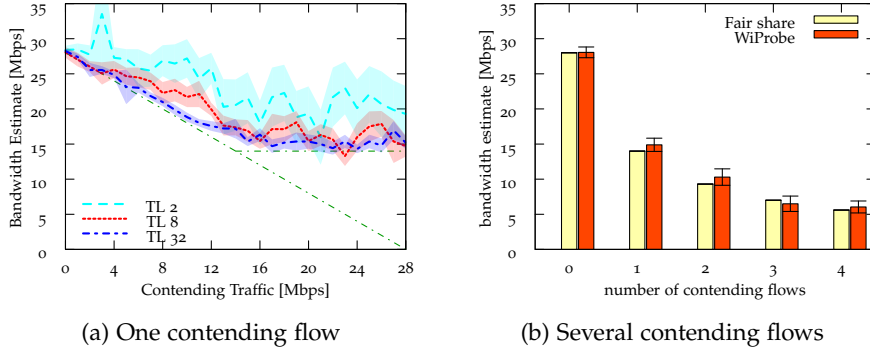
(a) One contending flow

(b) Several contending flows

Figure 24: Accuracy of WiProbe fair share probing with respect to different packet train length. We find that small packet trains yield incorrect estimates and a high variability whereas trains of medium length, i.e. $l \geq 8$, yield fairly good estimates and a small variation. Moreover, the estimates are quite accurate, also in cases of more than two stations.

Active Continuous Probing

First, we focus on the accuracy of unfiltered packet train probes. We estimate the fair share by sending probing traffic from a sender to a receiver. Concurrently, a second flow with a rate λ from increasing from 0 Mbps to 28 Mbps in 1 Mbps steps, contents for the same channel. This cross-traffic is generated using the iperf [5] traffic generator. Figure 24 shows the average of 25 fair share and bandwidth estimates with different probing train lengths gathered in approximately 1 second per rate for all cross-traffic rates. Furthermore, we show the corresponding confidence intervals at a confidence level of 0.95. As a reference, the true fair share of a new flow $f_s = max(C - \lambda, C/2)$ and the available bandwidth are plotted. As already stated in Section 6.6, we find that for medium-sized probing trains, i.e. $l \geq 8$, the fair share estimates are quite accurate. Shorter train lengths, however, result in samples with a higher variance, as indicated by the larger confidence intervals.

Figure 25 depicts the effect of train length and process noise on the ability to track a fair share discontinuity caused by Pareto cross-traffic. Based on the results plotted in Figures 23 and 24 we use a probing train length of $l = 8$ to detect the change of 14 Mbps for $T_s = 30$ s, i.e., a second station

(a) Influence of process noise on change tracking.

(b) Influence of process noise on convergence speed.

Figure 25: Measurement results of smoothed Kalman estimates from samples that are corrupted by measurements noise. The results restate the influence of process noise on change tracking and convergence speed, respectively, also for real systems with Pareto cross-traffic.

transmits its data with $r \geq 14$ Mbps for 30 seconds. Using the calculated, optimal process noise $\sigma_P^2 = 7.110^{-5}$, the jump is accurately tracked at the expense of a higher estimate noise. Nevertheless, the measurement noise is significantly lower when compared to the unfiltered case. Using a suboptimal value for σ_P^2, the filtered estimate variance is reduced; however, the estimate of the discontinuity is highly distorted.

The measurement results depicted in Figure 25 confirm our theoretical findings from Section 6.6. We were able to achieve a predetermined convergence time by selecting the process noise σ_P^2 according to equation (6.13). As expected, fast convergence times are associated with a larger estimate variance as indicated by the larger confidence intervals in Figure 25b.

Furthermore, we compared WiProbe to Pathchirp that also allows for continuous probing. Figure 26 depicts the mean bandwidth estimates over 25 experiments as well as corresponding confidence intervals for both tools. Regarding Pathchirp we used the default configuration and find that it underestimates the fair share significantly. This is in line with the results of Section 6.3. Moreover, Pathchirp detected the bandwidth change and followed the discontinuity with a convergence time of $T \approx 10$ s, see Figure 26a. We note, however, that, although it has several tuning knobs, it

(a) Pathchirp

(b) WiProbe

Figure 26: Continuous probing with WiProbe and Pathchirp. We find that both tools follow the bandwidth change. However, with a similar variance in the results, WiProbe provides estimates that are more accurate. Moreover, contrary to Pathchirp, WiProbe can be systematically parameterized to follow bandwidth changes over time with a predefined accuracy.

is not clear how to parameterize Pathchirp in order to detect bandwidth changes optimally.

For WiProbe we used a parameterization that achieves similar probing characteristics in terms of probing traffic intensity, probing frequency, and convergence time, as the default Pathchirp configuration. To this end, Figure 26b illustrated results for an optimal configuration with respect to the given discontinuity period of duration $T_s = 30$ s. We find that, according to the confidence intervals, both tools achieve comparable estimation variances. However, compared to the theoretic fair share baseline, WiProbe generally provides the more accurate results. Moreover, WiProbe can be tuned to detect shorter discontinuities optimally, whereas there is no such methodology to tune Pathchirp accordingly. Hence, it is not clear whether and how Pathchirp can track bandwidth changes over time that are smaller than $B/T \approx 1.4$.

Video Probing

The Kalman filter approach to available bandwidth estimation not only allows for active, but also for passive probing, e.g., by using a video data stream. To this end, every video frame on the application layer results

(a) video with CBR cross-traffic   (b) video with Pareto cross-traffic

Figure 27: Samples of the fair share show a large variability referred to as measurement noise. The noise reduces with $\sqrt{l}$. With this input the Kalman filter provides good estimates of the fair share after short convergence both for CBR as well as Pareto cross-traffic and for several contending stations.

in a packet burst on the network layer. Thus, with a video frame rate of 25 pictures per second, we obtain 25 bursts per second. The burst length corresponds to the video frame size. In the sequel, we perceive video bursts as probing trains. Using the gathered bandwidth information, a client-server video application can adapt its sending rate with respect to the available bandwidth. This approach has also been presented in [116], however, for wired networks with a focus on the video rate adaptation algorithm using scalable video coding [131], rather than the available bandwidth estimation.

To demonstrate the passive bandwidth estimation approach for wireless networks we applied the Video Lan Client (VLC) video player [9] to transmit a variable bit rate video tracefile[7] using UDP. In addition, the D-ITG traffic generator [1] is used at contending stations to generate CBR and Pareto cross-traffic with a shape parameter of 1.4 and changing intensity respectively. We used packet bursts caused by the video frames, approximately every 40 ms, to obtain a sample of the fair share. Samples obtained by bursts with a size $l \geq 8$ packets are fed into a Kalman filter that generates smoothed fair share estimates. In our experiments, we optimized the filter configuration for $M = 2$ and, with respect to IEEE 802.11g,

---

7 Some detailed information regarding the video trace file can be found in the Appendix on page 113.

$\Delta + L/C \approx 0.32$ ms such that the measurement noise has $\sigma_{g_d} \approx 0.49$ ms$/\sqrt{l}$. The Kalman filter weights each of the samples according to the corresponding measurement noise. Furthermore, we adjusted the Kalman filter to optimally track discontinuities of $T_s = 30$ ms, i.e. $\sigma_p^2 = 7.110^{-5}$ ms^2.

Figure 27 shows the individual samples derived from video frames of one measurement run. Additionally, it depicts smoothed fair share estimates from the Kalman filter averaged over 10 runs as well as the corresponding 0.95 confidence intervals. Despite the large variability of single fair share samples, the Kalman estimates follow the theoretical long-term fair share from equation (6.4) closely. Furthermore, the filter quickly detects changes in the underlying process, i.e. changes of the cross-traffic. This demonstrates the utility of our approach.

CONCLUSIONS AND FUTURE WORK

7.1 CONCLUSIONS

In this thesis we analyzed the short- and long-term fairness of the DCF using conditional probabilities of the number of inter-transmissions. The approach has been proven to be highly useful, facilitating significant closed-form results. Regarding i.i.d. countdown values, we showed a major impact of the type, but not the parameters of the distribution. We proved that long-term fairness improves with $\sqrt{l}$. Our findings are substantiated by an extensive measurement and simulation study. We modeled the DCF as emulating a fluid GPS scheduler yielding a fair average service rate that is subject to well-defined error terms. Based on the DCF clock, we derived a service curve that opens up significant options for performance analysis using the stochastic network calculus.

Moreover, we conducted an extensive measurement study of wireless bandwidth estimation in IEEE 802.11g WLAN testbeds. In contrast to wired links bandwidth estimates for wireless channels depend largely on the choice of packet sizes. We adapted the examined tools accordingly. We found that the FCFS assumption common in bandwidth estimation does not apply in case of wireless channels with contending traffic, where the distributed coordination function seeks to achieve a fair bandwidth allocation. We showed that the estimates of known iterative measurement methods can be related to the fair share of a new flow, which may deviate significantly from the available bandwidth. Our measurement results confirm this relation. A similar result was not established for direct probing.

Based on our findings, we developed our own fair share estimation method. We showed the benefit of employing Kalman filtering to improve estimates derived from continuous active probing of fair share using a

constant probing rate. Specifically, we showed how filter parameters should be chosen to fit a specific use case and calculated the effects of filtering on the estimate variance. Additionally, we derived the relationship between the filter convergence time and the process and measurement noise parameters. Furthermore, we evaluated the influence of train lengths on the estimate's variance for wired and wireless setups. We conclude that depending on the time scale and intensity of the cross-traffic variations of interest, it can be beneficial to sample cross-traffic frequently using short packet trains. Moreover, we proved that the estimation variance, that is inversely related to the smoothness of the estimates, decreases with $\sqrt{r_p}$.

We concluded our study showing a technique that estimates the fair rate from passive traffic measurements of a video application, e.g., for rate adaptation.

7.2   FUTURE WORK

Some open research questions remain from the findings presented in this thesis. We outline these questions in conjunction with possible applications.

As stated before, the fairness model assumes independent backoff values and, therefore, independent packet arrivals at the channel. This assumption is violated to some extent for $2 < M < \infty$. Thus, for a moderate number of hosts the model outcome differs slightly from the empirical data. However, it might be possible to find a correction term related to the auto-covariance of the packet arrivals that compensates this deviation.

The fairness model considers a basic service set and neglects hidden and exposed terminals as well as multi-hop scenarios. Hence, it would be interesting to investigate fairness in those cases. To this end, a stochastic network calculus approach, which allows the concatenation of systems, may reveal new insights also for end-to-end scenarios.

Regarding the upcoming IEEE 802.11n standard that incorporates new techniques such as MIMO, frame aggregation and service differentiation, a modified fairness model is needed.

Regarding the application of video-rate probing, a video application could benefit from the accurate bandwidth estimates by dynamically adapting the video rate according to the achievable bandwidth using, e.g., scalable video coding. Moreover, it would be interesting to investigate the relation of the process noise, that is needed for the Kalman filter parameterization, and the buffer size of the video application. To this end, one would need to concatenate traffic models of video sources with the Kalman filter approach.

Part II

APPENDIX

A

APPENDIX: PROOFS AND DERIVATIONS

PROOF 1: RANDOM SUM OF RANDOM VARIABLES

This proof corresponds to the derivation of packet departure time in IEEE 802.11 WLANs including retransmissions on page 63.

In the following, we show that the expected value of a random sum of i.i.d. random variables can be calculated by the expected values of each random variable. That is,

$$E\left[\sum_{j=1}^{X} Y_j\right] = E\left[X\right] E\left[Y\right]$$

where X and Y are independent random variables which are negative binomial and exponentially distributed, respectively.

Proof The proof uses moment generating functions. Given that X is negative binomial distributed with parameters r and p, we find the moment generating function by

$$M_X(\theta) = \left(\frac{1-p}{1-p\,e^\theta}\right)^r$$

Moreover, since Y is exponentially distributed with parameter λ, we find

$$M_Y(\theta) = \frac{\lambda}{\lambda - \theta}$$

If X_1, X_2 ... is a sequence of i.i.d. random variables, and $Y \geq 0$ is a random variable which is independent of X, then the sum $S = X_1 + X_2 + \cdots + X_Y$ has a moment generating function given by $M_S(\theta) = M_X\left(ln\left(M_Y(\theta)\right)\right)$,

see, e.g., [62, 148]. Thus, we calculate the moment generation function of the random sum by

$$M_S(\theta) = \left(\frac{1-p}{1 - p \, e^{ln\left(\frac{\lambda}{\lambda - \theta}\right)}} \right)^r$$

Next, we derive the first moment by differentiating the moment generating functions at $\theta = 0$. That is

$$\frac{d}{d\theta} M_X(0) = \frac{rp}{1-p} \qquad\qquad = E[X]$$

$$\frac{d}{d\theta} M_Y(0) = \frac{1}{\lambda} \qquad\qquad = E[Y]$$

$$\frac{d}{d\theta} M_S(0) = \frac{rp}{\lambda(1-p)} \quad = E\left[\sum_{j=1}^{X} Y_i \right]$$

Now, it is evident, that $E[X]\,E[Y] = E\left[\sum_{j=1}^{X} Y_i\right]$, i.e. $\frac{1}{\lambda} \frac{rp}{1-p} = \frac{rp}{\lambda(1-p)}$. To this end, we find that $E[\phi_i(n)] = E[\sum_{j=1}^{1+R} b_i(n) - (1 + E[R])\mu] = 0$ as stated in the text. ∎

A general proof for the random sum of random variables, that holds for any distribution with a known moment generating function, can be found, e.g., in [62].

PROOF 2: $q_1 < 1$

This proof corresponds to the calculation of violation probabilities and sample path bounds on page 67

We use Boole's inequality to sum the violation probabilities ε_i over all packets l to derive a corresponding sample path bound. To this end, we need to prove that the sum of the probabilities of the individual events stays finite. This can be done using the geometric sum. However, this requires the probability of any single event to be less than one, i.e. $q_1 < 1$.

Proof From equation (5.8) we find

$$\varepsilon_1(l') \leq \left(\frac{\vartheta}{\mu} e^{-\frac{\mu+\vartheta}{\mu}} \right)^{l'} = q_1^{l'}$$

which yields

$$q_1 = \frac{\vartheta}{\mu} e^{-\frac{\mu+\vartheta}{\mu}}.$$

Substituting ϑ/μ by a yields

$$q_1 = a \, e^{-(a+1)}$$

Now, it is evident that q_1 is negative, i.e $q_1 < 0$ if $a = \vartheta/\mu$, is negative. Moreover, we differentiate q_1 two times to find its maximum value

$$\frac{d}{da} q_1 = e^{-(a+1)} - a \, e^{-(a+1)}$$
$$\frac{d^2}{d^2 a} q_1 = -2 \, e^{-(a+1)} - a \, e^{-(a+1)}$$

Solving the first derivative $\frac{d}{d\rho} q_1 = 0$ for a we find that q_1 has an extreme value at $(1, e^{-2})$. Moreover, this has to be a global maximum since $\frac{d^2}{d^2\rho} q_1 \big|_{a=1} < 0$ and $lim_{a\to\infty} q_1 = 0$. Hence, $q_1 < 1$ holds generally. ∎

PROOF 3: $q_3 < 1$

This proof corresponds to the calculation of violation probabilities and sample path bounds on page 68

For reasons similar to proof 1, we need to show that q_2 and q_3, respectively, are less than 1. In the sequel we outline the proof for $q_3 < 1$, since the notation is slightly simpler. However, the proof for $q_2 < 1$ follows accordingly.

Proof We have from equation (5.10) that

$$\varepsilon_3(l) \leq \left(\frac{p(1-p)^\rho(1+\rho)^{1+\rho}}{\rho^\rho} \right)^l = q_3^l$$

which yields

$$q_3 = \frac{p(1-p)^\rho(1+\rho)^{1+\rho}}{\rho^\rho}.$$

To show that $q_3 < 1$ we, again, derive $\frac{d}{d\rho} q_3$

$$\frac{d}{d\rho} q_3 = \frac{p(1-p)^\rho (1+\rho)^{1+\rho}}{\rho^\rho} \left[\ln(1-p) + \ln(1+\rho) - \ln\rho) \right]$$

and solve $\frac{d}{d\rho} q_3 = 0$ for ρ to find an extreme value at $\rho = (1-p)/p$. Since the constant $p = 1/M$ in $(0,1]$ reflects the channel access probability as a function of number of contending stations M, we find that $(1-p)/p = M - 1$. To this end, we can calculate the corresponding extreme value by

$$q_3 \big|_{\rho = \frac{1-p}{p}} = \frac{p(1-p)^{\frac{1-p}{p}} (1 + \frac{1-p}{p})^{1+\frac{1-p}{p}}}{\frac{1-p}{p}^{\frac{1-p}{p}}}$$

$$= \frac{(\frac{M-1}{M})^{M-1} M^{M-1}}{(M-1)^{M-1}}$$

After some simple reordering we find

$$q_3 \big|_{\rho = \frac{1-p}{p}} = 1.$$

To prove that $(\frac{1-p}{p}, 1)$ is a maximum we need to show $\frac{d^2}{d^2\rho} q_3 \big|_{\rho = \frac{1-p}{p}} < 0$. For this purpose, we calculate the derivative by

$$\frac{d^2}{d^2\rho} q_3 = \frac{p(1-p)^\rho (1+\rho)^{1+\rho}}{\rho^\rho} \left[\frac{1}{(1+\rho)} - \frac{1}{\rho} + \ln(1-p)^2 \right.$$

$$+ 2\ln(1-p)(\ln(1+\rho)+1) + (\ln(1+\rho)+1)^2$$

$$- 2\ln(1-p)(\ln(\rho)+1) + (\ln(\rho)+1)^2$$

$$\left. - 2(\ln(1+\rho)+1)(\ln(\rho)+1) \right]$$

Solving the above equation for $M = 2$ stations, i.e. $p = 1/2$, yields $\frac{d^2}{d^2\rho}q_3\big|_{\rho=\frac{1-p}{p}} = -0.5$. Hence, $(\frac{1-p}{p}, 1)$ is a maximum value.

Moreover, to prove this a global maximum, we explore the tail ends of q_2. We find that $\lim_{\rho\to-\infty} q_3 = -\infty$ and $\lim_{\rho\to\infty} = 0$ respectively. Thus, for every $p \in (0, 1]$ and for every $\rho \lessgtr \frac{1-p}{p}$, the function q_3 is less than 1. ∎

DERIVATION 1: PROCESS NOISE

This derivation corresponds to the calculation the process noise of a random bandwidth process on page 89

In the sequel, we show that for the process noise $\sigma_P^2 = g_B^2/n_p$ holds generally. We use that $g_B = L/B$, where L is the packet size and $B = B_{max} - B_{min}$ is the maximum bandwidth change. Furthermore, n_p is the number of samples used to detect a discontinuity of a given length. From the definition of the unbiased sample variance, it follows that

$$Var(x) = \frac{1}{N-1} \sum_{i=1}^{N} \left(x_i - \frac{1}{N} \sum_{j=1}^{N} x_j \right)^2 .$$

Since we are interested in the maximum difference of bandwidth changes that has to be detected in $N = n_p$ steps, we can assume that $x_1 = g_B$ while $x_i = 0$ for $i > 1$. Thus, the above equation reduces to

$$Var(x) = \frac{1}{N-1} \left(\frac{(Nx_1 - x_1)^2}{N^2} + \frac{(N-1)x_1^2}{N^2} \right) .$$

After some reordering it follows that

$$Var(x) = \frac{x_1^2}{N} = \frac{g_B^2}{n_p} = \sigma_P^2 .$$

DERIVATION 2: OPTIMAL TRAIN LENGTH

This derivation corresponds to the calculation the optimal packet train length for continuous bandwidth estimation on page 93

To derive the optimal train length, we have to minimize equation (6.14). Thus, we have to differentiate equation (6.14) with respect to l, i.e.

$$\frac{d}{dl}\sigma^2_{\hat{x}(n)} = \frac{G}{2-G}\left(\frac{g_B^2 L}{T_s r_p} - \frac{p^2 p_c + (1-p)(1-p_c)p^2 l^2}{(1-p_c)^2}(\Delta + L/C)^2\right)$$

Solving $\frac{d}{dl}\sigma^2_{\hat{x}(n)} = 0$ for l yields the packet train length that minimized equation (6.14). Thus, the optimal train length follows by the positive result of

$$l = \frac{\sqrt{L\,T_s\,r_p\,(p^2 p_c^2 + p p_c - p - p_c + 1)}}{L\,(1-p_c)\,p\,g_B}\,(\Delta + L/C).$$

APPENDIX: VIDEO

In the following, we present some information and statistics regarding the video trace file used in our experiments. We used serveral episodes of the film Spiderman. The film was encoded as single layer MPEG-2 with a group-of-picture size of 12 video frames. The film has an overal size of approx. 1766 Mbyte and lasts for 5036 s, i.e. approx. 84 min. This results in an average probing rate r_p of approx. 2.8 Mbps.
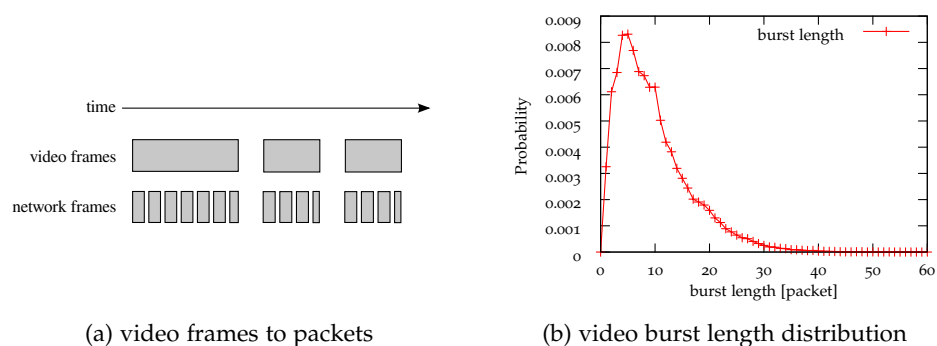

| (a) video frames to packets | (b) video burst length distribution |

Figure 28: The burst length distribution of the video trace file used in our experiments. As illustrated, any video frame results in a packet burst on the network layer. The average burst size of our video is 9 packets per burst.

As illustrated in Figure 28a, any video frame on the application layer, results in a packet burst of a specific length on the network layer. The average bust length l is given by 9 packets/burst with a packet size L of 1442 byte. The distribution of bust length is depicted in Figure 28b. Moreover, the minimum video frame size is 58 byte, i.e. one 64 byte packet on the network layer. The maximum video frame size is 87 kbyte which results in a bust length of 60 packets.

ACK	Acknowledgement
AIMD	Additive increase, multiplicative decrease
AR	Auto-regressive
ARQ	Automatic repeat request
ATM	Asynchronous transfer mode
AvBw	Available bandwidth
BEB	Binary exponential backoff
BSS	Basic service set
Cor.	Corollary
CSMA	Carrier sense multiple access
CTS	Clear to send
DCF	Distributed coordination function
DIFS	Distributed (coordination function) inter-frame space
DLS	Direct link setup
DRR	Deficit round robin
DSSS	Direct sequence spread spectrum
Eq.	Equation
ERP	Extended rate physical layer
EWMA	Exponentially weighted moving average
FCFS	First come first serve
FIFO	First in first out
Fig.	Figure
Gbps	Gigabit per second
GoP	Group of pictures
GPS	Generalized processor sharing
GRC	Guaranteed rate clock
GSM	Global System for Mobile Communications
HR	High rate
i.i.d	independent and identically distributed
IEEE	Institute of Electrical and Electronics Engineers
ISM	Industrial, scientific, and medical
LAN	Local Area Network
MAC	Medium access control
MACA	Multiple access with collision avoidance

MACAW	Multiple access with collision avoidance for wireless
Mbps	Mega bit per second
MILD	Multiplicative increase, linear decrease
MIMO	Multiple input, multiple output
MSE	Mean squared error
OFDM	Orthogonal frequency division multiplexing
PCF	Point coordination function
PHY	Physical (layer)
PIFS	Point (coordination function) inter-frame space
PLCP	Physical layer convergence procedure
QoS	Quality of service
RLC	Radio link control
RTS	Request to send
SIFS	Short inter-frame space
SISO	Single input, single output
TCP	Transport control protocol
Th.	Theorem
TXOP	Transmit opportunity
UDP	User datagram protocol
WFQ	Weighted fair queuing
WLAN	Wireless Local Area Network

$(x)^+$	Positive part $max(0,x)$ of a number x
$A(t)$	Arrival process
$AvBw$	Available bandwidth
C	Channel capacity
$D(t)$	Departure process
$E[X]$	Expectation value of a random variable X
G	Goodput, i.e. achievable throughput
J	Jain's fairness index
L	Packet size
M	Number of wireless stations
$P[X]$	probability of random variable X
$S(t)$	Service curve function
$Var[X]$	Variance of a random variable X
$a(n)$	Arrival timestamp of packet n
b	i.i.d. random backoff variable
c	A constant
$d(n)$	Departure timestamp of packet n
f	Fair bandwidth share
g_d	Average gap of departing packets
g_i	Input gap, related to the input rate
g_o	Output gap, related to the output rate
i, j, n, m	general indexes
k	Inter-transmitted packets
l	Packet burst length
p	channel access probability
p_a	transmission attempt probability
p_c	collision probability
r	Data rate
r_i	Input rate to a system
r_o	Output rate to a system
u	Collision window of IEEE 802.3
w	Contention window of the DCF
$M_X(\theta$	Moment generating function of X
$N(\mu,\sigma)$	Normal random variable with expected value μ and variance σ
$S^\varepsilon(n)$	Stochastic service curve function

$U(\tau, t)$	Channel utilization in the interval $[\tau, t]$
$\hat{M}$	Tagged wireless station
ι	The lag of an autocorrelated process
κ	Packet arrival process
$s^\varepsilon(n)$	Stochastic service timestamp of packet n
t, τ	Time indexes
$*$	Max-plus convolution operator
$\otimes$	Min-plus convolution operator
ε	Percentile that describes the violation probability
λ	Cross traffic rate
ϕ_i	Weight assigned to flow i in a weighted fair queuing system
$\blacksquare$	Halmos symbol to end of a proof

[1] D-ITG: Internet Traffic Generator. `http://www.grid.unina.it/software/ITG`.

[2] DD-WRT: Linux Wireless Router. `http://www.dd-wrt.com`.

[3] Gateworks Laguna Network Platform. `http://www.gateworks.com`.

[4] INETMANET Framework for OMNEST/OMNeT++ 4.x. `http://github.com/inetmanet`.

[5] Iperf: TCP/UDP Bandwidth Measurement Tool. `http://dast.nlanr.net/Projects/Iperf`.

[6] The Linux Kernel Archive. `http://www.kernel.org`.

[7] OMNeT++: Network Simulator. `http://www.omnetpp.org`.

[8] OpenWrt: Linux Wireless Router. `http://www.openwrt.org`.

[9] VLC Video Lan Client. `http://www.videolan.org`.

[10] A Hybrid Wireless MAC Protocol Supporting Asynchronous and Synchronous MSDU Delivery Services, September 1992. IEEE 802.11 Working Group paper 802.11/91-92.

[11] IEEE Standard for Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Spedifications, June 2007. Std. 802.11-2007.

[12] IEEE Standard for Carrier Sense Multiple Access with Collision Detection (CMSA/CD) Access Method and Physical Layer Specifications, December 2008. Std. 802.3-2008.

[13] O. Abboud, T. Zinner, K. Pussep, S. A.-S., and R. Steinmetz. On the Impact of Quality Adaptation in SVC-based P2P Video-on-Demand Systems. In *ACM Multimedia Systems*, pages 177–197, February 2011.

[14] N. Abramson. The ALOHA system - another alternative for computer communications. In *Proc. of Fall Joint Computing Conference*, pages 281–285, May 1970.

[15] N. Abramson. The Throughput of Packet Broadcasting Channels. *IEEE Trans. Commun.*, 25(1):117–128, January 1977.

[16] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-Level Measurements from an 802.11b Mesh Network. *SIGCOMM Comput. Commun. Rev.*, 34:121–132, August 2004.

[17] A. Amambra, K. M. Hou, and J.-P. Chanet. Evaluation of the Performance of the SLoPS: Available Bandwidth Estimation Technique in IEEE 802.11b Wireless Networks. In *Proc. of IFIP NTMS*, pages 123–132, May 2007.

[18] A. Antoniou. *Digital Filters: Analysis, Design, and Applications*. McGraw-Hill, 1993.

[19] G. Armitage. An Experimental Estimation of Latency Sensitivity in Multiplayer Quake 3. In *Proc. of IEEE ICON*, pages 137–141, September 2003.

[20] A. V. Babu and L. Jacob. Fairness analysis of IEEE 802.11 multirate wireless LANs. *IEEE Trans. Veh. Technol.*, pages 3073–3088, September 2007.

[21] F. Baccelli, G. Cohen, and G. Oslder. *Synchronization and Linearity: An Algebra for Discrete Event Systems*. John Wiley & Sons Inc., 1992.

[22] A. Balachandran, G. M. Voelker, P. Bahl, and P. V. Rangan. Characterizing User Behavior and Network Performance in a Public Wireless LAN. *SIGMETRICS Perform. Eval. Rev.*, 30:195–205, June 2002.

[23] A. Banchs and X. Perez. Distributed Weighted Fair Queuing in 802.11 Wireless LAN. In *Proc. of IEEE ICC*, pages 3121–3127, May 2002.

[24] C. L. Barrett, M. V. Marathe, D. C. Engelhart, and A. Sivasubramaniam. Analyzing the Short-Term Fairness of IEEE 802.11 in Wireless Multi-hop Radio Networks. In *Proc. of IEEE MASCOTS*, pages 137–144, October 2002.

[25] J. Beran. *Statistics for Long-Memory Processes*. Chapman & Hall/CRC, 1994.

[26] J. Beran, R. Sherman, M. S. Taqqu, and W. Willinger. Long-Range Dependence in Variable-Bit-Rate Video Traffic. *IEEE Trans. Commun.*, 43(234):1566–1579, February 1995.

[27] G. Berger-Sabbatel, A. Duda, O. Gaudoin, M. Heusse, and F. Rousseau. Fairness and its impact on delay in 802.11 networks. In *Proc. of IEEE GLOBECOM*, pages 2967–2973, November 2004.

[28] G. Berger-Sabbatel, A. Duda, O. Gaudoin, M. Heusse, and F. Rousseau. Short-term Fairness of 802.11 Networks with Several Hosts. In *Proc. of IFIP MWCN*, pages 263–274, October 2004.

[29] D. Bertsekas and R. Gallager. *Data Networks*. Pearson Prentice Hall, 1979.

[30] V. Bharghavan, A. Demers, S. Shenker, and L. Zhang. MACAW: A Media Access Protocol for Wireless LAN's. In *Proc. of ACM SIGCOMM*, pages 212–225, October 1994.

[31] G. Bianchi. Performance Analysis of the IEEE 802.11 Distributed Coordination Function. *IEEE J. Select. Areas Commun.*, 18(3):535–547, March 2000.

[32] G. Bolch, S. Greiner, H. de Meer, and K. S. Trivedi. *Queueing Networks and Markov Chains*. John Wiley & Sons Inc., 2006.

[33] R. R. Boorstyn, A. Burchard, J. Liebeherr, and C. Oottamakorn. Statistical Service Assurances for Traffic Scheduling Algorithms. *IEEE Journal on Selected Areas in Communications*, 18(12):2651–2664, December 2000.

[34] A. Botta, A. Pescape, and G. Ventre. On the Performance of Bandwidth Estimation Tools. In *Proc. of IEEE Systems Communications*, pages 287–292, August 2005.

[35] Z. Bozakov and M. Bredel. SSHLauncher - A Tool for Experiment Automation in Distributed Environments. Technical report, TU Darmstadt, July 2008.

[36] Z. Bozakov and M. Bredel. Online Estimation of Available Bandwidth and Fair Share Using Kalman Filtering. In *Proc. of Networking*, pages 548–561, May 2009.

[37] R. Braden, D. Clark, and S. Shenker. RFC1633 - Integrated Services in the Internet Architecture: an Overview. `http://www.ietf.org/rfc/rfc1633.txt`, June 1994.

[38] M. Bredel and M. Fidler. A Measurement Study of Bandwidth Estimation in IEEE 802.11g Wireless LANs using the DCF. In *Proc. of IFIP Networking*, pages 314–325, May 2008.

[39] M. Bredel and M. Fidler. A Measurement Study regarding Quality of Service and its Impact on Multiplayer Online Games. In *Proc. of NetGames*, pages 1–6, November 2010.

[40] R. Bringhurst. *The Elements of Typographic Style*. Version 2.5. Hartley & Marks, Publishers, 2002.

[41] F. Calì, M. Conti, and E. Gregori. Dynamic Tuning of the IEEE 802.11 Protocol to Achieve a Theoretical Throughput Limit. *IEEE/ACM Trans. Networking*, 8(6):785–799, December 2000.

[42] M. Carvalho and J. J. Garcia-Luna-Aceves. Delay Analysis of IEEE 802.11 in Single-Hop Networks. In *Proc. of ICNP*, pages 146–155, November 2003.

[43] C. S. Chang. *Performance Guarantees in Communication Networks*. John Wiley & Sons Inc., 2000.

[44] J. I. Choi, M. Jain, K. Srinivasan, P. Levis, and S. Katti. Achieving single channel, full duplex wireless communication. In *Proc. of ACM MobiCom*, pages 1–12, September 2010.

[45] C. K. Chui and G. Chen. *Kalman Filtering with Real-Time Applications*. Springer Press, 2009.

[46] F. Ciucu. *Scaling Properties in the Stochastic Network Calculus*. PhD thesis, University of Virginia, August 2007.

[47] F. Ciucu, A. Burchard, and J. Liebeherr. Scaling Properties of Statistical End-to-End Bounds in the Network Calculus. *IEEE/ACM Trans. Networking*, 14(6):2300–2312, June 2006.

[48] S. Cocorada. An IEEE 802.11g Simulation Model with Extended Debug Capabilities. In *Proc. of SimuTools*, pages 1–3, March 2008.

[49] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. John Wiley & Sons Inc., 1991.

[50] M. E. Crovella and A. Bestavros. Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes. In *Proc. of ACM SIGMETRICS*, pages 160–169, May 1996.

[51] R. L. Cruz. A Calculus for Network Delay. *IEEE Trans. Inform. Theory*, 37(1):114–141, January 1991.

[52] R. L. Cruz and C. M. Okino. Service Guarantees for Window Flow Control. In *Proc. Allerton Conf. on Comm., Cont. & Comp*, pages 1–12, October 1996.

[53] A. Demers, S. Keshav, and S. Shenker. Analysis and Simulation of a Fair Queueing Algorithm. In *Proc. of ACM SIGCOMM*, pages 1–12, October 1989.

[54] A. Ekelin, M. Nilsson, E. Hartikainen, A. Johnsson, J.-E. Mangs, B. Melander, and M. Björkman. Real-time Measurement of End-To-End Available Bandwidth Using Kalman Filtering. In *Proc. IEEE/IFIP NOMS*, pages 5–17, April 2006.

[55] A. K. Erlang. The Theory of Probabilities and Telephone Conversations. *Nyt Tidsskrift for Matematik*, 20(1):33–39, 1909.

[56] A. K. Erlang. Solution of some Problems in the Theory of Probabilities of Significance in Automatic Telephone Exchanges. *Elektrotkeknikeren*, 13(1):5–22, 1917.

[57] M. Fidler. A Network Calculus Approach to Probabilistic Quality of Service Analysis of Fading Channels. In *Proc. of IEEE GLOBECOM*, pages 1–6, November 2006.

[58] M. Fidler. An End-to-End Probabilistic Network Calculus with Moment Generating Functions. In *Proc. of IWQoS*, pages 261–270, June 2006.

[59] M. Fidler. A Survey of deterministic and stochastic service curve models in the network calculus. *IEEE Communications Surveys Tutorials*, 12(1):59–86, March 2010.

[60] M. Garetto, J. Shi, and E. W. Knightly. Modeling Media Access in Embedded Two-Flow Topologies of Multi-hop Wireless Networks. In *Proc. of ACM MobiCom*, pages 200–214, September 2005.

[61] P. Goyal, S. S. Lam, and H. M. Vin. Determining End-to-End Delay Bounds in Heterogeneous Networks. *Multimedia Systems*, 5(3):157–163, May 1997.

[62] G. Grimmett and D. Stirzaker. *Probability and Random Processes*. Oxford University Press, 1991.

[63] J. P. Hespana. *Linear System Theory*. Princeton University Press, 2009.

[64] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda. Performance anomaly of 802.11b. In *Proc. of IEEE INFOCOM*, pages 836–843, March 2003.

[65] C. Hoon, V. Misra, and D. Rubenstein. Fairness and Physical Layer Capture in Random Access Networks. In *Proc. of SECON*, pages 381–390, June 2007.

[66] N. Hu. *Network Monitoring and Diagnosis Based on Available Bandwidth Measurement*. PhD thesis, Carnegie Mellon University, May 2006.

[67] N. Hu and P. Steenkiste. Evaluation and Characterization of Available Bandwidth Probing Techniques. In *IEEE J. Select. Areas Commun.*, pages 879–894, August 2003.

[68] iSuppli Market Research. Wi-Fi Chipset Shipments to Double in 2011, February 2011. `http://www.isuppli.com` [Online; accessed 25-February-2011].

[69] J. R. Jackson. Networks of Waiting Lines. *Operations Research*, 5(5): 518–521, August 1957.

[70] V. Jacobson. Congestion Avoidance and Control. In *Proc. of ACM SIGCOMM*, pages 273–288, August 1988.

[71] M. Jain and C. Dovrolis. End-to-End Available Bandwidth: Measurement Methodology, Dynamics, and Relation with TCP Throughput. In *IEEE/ACM Trans. Networking*, pages 537–583, August 2003.

[72] M. Jain and C. Dovrolis. Ten Fallacies and Pitfalls on End-To-End Available Bandwidth Estimation. In *Proc. of ACM IMC*, pages 272–277, October 2004.

[73] M. Jain and C. Dovrolis. Path selection using available bandwidth estimation in overlay-based video streaming. *Computer Networks*, 52 (12):2411–2418, August 2008.

[74] R. Jain. *The Art of Computer Performance Analysis*. John Wiley & Sons Inc., 1991.

[75] R. Jain, D.-M. Chiu, and W. R. Hawe. A Quantitative Measure of Fairness and Discrimination for Resource Allocation in Shared Computer System. Technical Report DEC-TR-301, Digital Equipment, September 1984.

[76] Y. Jian and S. Chen. Can CSMA/CA networks Made be Fair? In *Proc. ACM MobiCom*, pages 137–144, September 2008.

[77] L. B. Jiang and S. C. Liew. Proportional fairness in wireless LANs and ad hoc networks. *IEEE WCNC*, pages 1551–1556, March 2005.

[78] Y. Jiang. Relationship between Guaranteed Rate Server and Latency Rate Server. *Computer Networks*, 43(3):307–315, October 2003.

[79] Y. Jiang. A Basic Stochastic Network Calculus. In *Proc. ACM SIG-COMM*, pages 123–134, October 2006.

[80] Y. Jiang and Y. Liu. *Stochastic Network Calculus*. Springer Press, 2008.

[81] A. Johnsson, B. Melander, and M. Björkman. Diettopp: A First Implementation and Evaluation of a Simplified Bandwidth Measurement Method. In *Proc. of SNCNW*, November 2004.

[82] A. Johnsson, B. Melander, and Ma. Björkman. Bandwidth Measurement in Wireless Networks. In *Proc. of Med-Hoc-Net*, June 2005.

[83] A. Johnsson, B. Melander, and Ma. Björkman. An Analysis of Active End-to-End Bandwidth Measurements in Wireless Networks. In *Proc. of IEEE/IFIP E3EMON*, pages 74–81, April 2006.

[84] T. Joshi, A. Mukherjee, Y. Yoo, and D.-P. Agrawal. Airtime Fairness for IEEE 802.11 Multirate Networks. *IEEE Trans. Mobile Computing*, pages 513–527, April 2008.

[85] J. Jun and M. L. Sichitiu. Fairness and QoS in Multihop Wireless Networks. In *Proc. of IEEE VTC*, pages 2936–2940, October 2003.

[86] P. Karn. MACA a New Channel Access Method for Packet Radio. In *Proc. of the Computer Networking Conference*, pages 134–140, June 1990.

[87] F. P. Kelly. Stochastic Models for Computer Communication Systems. *Journal of the Royal Statistic Society: Series B*, 47(3):379–395, 1985.

[88] F. P. Kelly. Notes on Effective Bandwidths. *Royal Statistical Society Lecture Notes*, 4:141–168, September 1996.

[89] S. Keshav. A Control-Theoretic Approach to Flow Control. In *Proc. of ACM SIGCOMM*, pages 3–15, September 1991.

[90] M. A. Y. Khan and D. Veitch. Isolating Physical PER for Smart Rate Selection in 802.11. In *Proc. of IEEE INFOCOM*, pages 1080–1088, April 2009.

[91] L. Kleinrock. *Message Delay in Communication Nets with Storage*. PhD thesis, Massachusetts Institute of Technology, December 1962.

[92] L. Kleinrock. *Queueing Systems*, volume I. John Wiley & Sons Inc., 1975.

[93] C. E. Koksal, H. Kassab, and H. Balakrishnan. An Analysis of Short-Term Fairness in Wireless Media Access Protocols. In *Proc. of ACM SIGMETRICS*, pages 118–119, June 2000.

[94] V. Konda and J. Kaur. RAPID: Shrinking the Congestion-Control Timescale. In *INFOCOM 2009, IEEE*, pages 1–9, April 2009.

[95] A. Kumar, D. Manjunath, and J. Kuri. *Communication Networking: An Analytical Approach*. Morgan Kaufmann, 2004.

[96] B.-J. Kwak, N.-O. Song, and L. E. Miller. Performance analysis of exponential backoff. *IEEE/ACM Trans. Networking*, 13(2):343–355, April 2005.

[97] Y. Kwon, Y. Fang, and H. Latchman. A novel MAC protocol with fast collision resolution for wireless LANs. In *Proc. of IEEE INFOCOM*, pages 853–862, April 2003.

[98] K. Lakshminarayanan, V. N. Padmanabhan, and J. Padhye. Bandwidth Estimation in Broadband Access Networks. In *Proc. of ACM IMC*, pages 314–321, October 2004.

[99] A. M. Law and W. D. Kelton. *Simulation Modeling and Analysis*. McGraw-Hill, 1991.

[100] J.-Y. Le Boudec. Application of Network Calculus to Guaranteed Service Networks. *IEEE Trans. Inform. Theory*, 44(3):1087–1096, May 1998.

[101] J.-Y. Le Boudec and P. Thiran. *Network Calculus A Theory of Deterministic Queuing Systems for the Internet*. Springer Press, 2001.

[102] H. K. Lee, V. Hall, K. H. Yum, K. I. Kim, and E. J. Kim. Bandwidth Estimation in Wireless LANs for Multimedia Streaming Services. In *Proc. of IEEE ICME*, pages 1181–1184, July 2006.

[103] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson. On the Self-Similar Nature of Ethernet Traffic. *IEEE/ACM Trans. Networking*, 2(1):1–15, February 1994.

[104] C. Li, H. Che, and S. Li. A Wireless Channel Capacity Model for Quality of Service. *IEEE Trans. Wireless Commun.*, 6(1):356–366, January 2007.

[105] M. Li, M. Claypool, and R. Kinicki. WBest: A Bandwidth Estimation Tool for Multimedia Streaming Application over IEEE 802.11 Wireless Networks. Technical Report WPI-CS-TR-06-14, Computer Science Department, Worcester Polytechnic Institute, March 2006.

[106] Z. Li, S. Nandi, and A. K. Gupta. Modeling the short-term unfairness of IEEE 802.11 in presence of hidden terminals. *Performance Evaluation*, 63(4):441–462, May 2006.

[107] J. Liebeherr, M. Fidler, and S. Valaee. A Min-Plus System Interpretation of Bandwidth Estimation. In *Proc. of IEEE INFOCOM*, pages 1127–1135, May 2007.

[108] X. Liu, K. Ravindran, and D. Loguinov. A Queuing-Theoretic Foundation of Available Bandwidth Estimation: Single-Hop Analysis. *IEEE/ACM Trans. Networking*, 15(2):918–931, August 2007.

[109] X. Liu, K. Ravindran, and D. Loguinov. A Stochastic Foundation of Available Bandwidth Estimation: Multi-hop Analysis. *IEEE/ACM Trans. Networking*, 16(2):130–143, April 2008.

[110] P. Loiseau, P. Goncandalves, G. Dewaele, P. Borgnat, P. Abry, and P. V.-B. Primet. Investigating Self-Similarity and Heavy-Tailed Distributions on a Large-Scale Experimental Facility. *IEEE/ACM Trans. Networking*, 18(4):1261–1274, August 2010.

[111] R. Lübben, M. Fidler, and J. Liebeherr. A Foundation for Stochastic Bandwidth Estimation of Networks with Random Service. In *Proc. of IEEE INFOCOM*, pages 1–9, March 2011.

[112] S. Machiraju, D. Veitch, F. Baccelli, and J. Bolot. Adding Definition to Active Probing. *ACM SIGCOMM Computer Communication Review*, 37 (2):19–28, 2007.

[113] K. Mahmood, A. Rizk, and Y. Jiang. On the Flow-Level Delay of a Spatial Multiplexing MIMO Wireless Channel. In *Proc. of IEEE ICC*, pages 53–59, June 2011.

[114] B. Melander, M. Björkman, and P. Gunningberg. Regression-Based Available Bandwidth Measurements. In *Proc. of SPECTS*, pages 38–63, July 2002.

[115] T. Nandagopal, T.-E. Kim, X. Gao, and V. Bharghavan. Achieving MAC Layer Fairness in Wireless Packet Networks. In *Proc. of ACM MobiCom*, pages 87–98, September 2000.

[116] D. T. Nguyen and J. Ostermann. Congestion Control for Scalable Video Streaming Using the Scalability Extension of H.264/AVC. *IEEE Journal of Selected Topics in Signal Processing*, 1(2):246–253, August 2007.

[117] I. Norros. A Storage Model with Self-Similar Input. *Queueing Systems*, 16(3):387–396, September 1994.

[118] I. Norros. On the Use of Fractional Brownian Motion in the Theory of Connectionless Networks. *IEEE J. Sel. Areas Commun*, 13(6):953–962, August 1995.

[119] A. Papoulis and S. U. Pillai. *Probability, Random Variables, and Stochastic Processes*. McGraw-Hill, 2002.

[120] A. K. Parekh and R. G. Gallager. A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks. *IEEE/ACM Trans. Networking*, 1(3):344–357, June 1993.

[121] V. Paxson and S. Floyd. Wide area traffic: the failure of Poisson modeling. *IEEE/ACM Trans. Networking*, 3(3):226–244, June 1995.

[122] E. Perahia and R. Stacey. *Next Generation Wireless LANs - Throughput, Robustness, and Reliability in 802.11n*. Cambridge University Press, 2008.

[123] C. L. Phillips, J. M. Parr, and E. A. Riskin. *Signals, Systems and Transforms*. Pearson Prentice Hall, 2007.

[124] M. Portoles-Comeras, A. Cabellos-Aparicio, J. Mangues-Bafalluy, A. Banchs, and J. Domingo-Pascual. Impact of Transient CSMA/CA Access Delays on Active Bandwidth Measurements. In *Proc. of ACM IMC*, pages 397–409, November 2009.

[125] V. J. Ribeiro, R. H. Riedi, R. G. Baraniuk, J. Navratil, and L. Cottrell. Pathchirp: Efficient Available Bandwidth Estimation for Network Paths. In *Proc. of PAM*, April 2003.

[126] L. G. Roberts. ALOHA Packet System With and Without Slots and Capture. *SIGCOMM Comput. Commun. Rev.*, 5:28–42, April 1975.

[127] S. Ross. *A First Course in Probability*. Pearson Prentice Hall, 2006.

[128] M. Roughan. Fundamental Bounds on the Accuracy of Network Performance Measurements. *SIGMETRICS Perform. Eval. Rev.*, 33: 253–264, June 2005.

[129] T. Sakurai and H. L. Vu. MAC Access Delay of IEEE 802.11 DCF. *IEEE Trans. Wireless Commun.*, 6(5):1702–1710, May 2007.

[130] C. Sarr, C. Chaudet, G. Chelius, and I. G. Lassous. Improving Accuracy in Available Bandwidth Estimation for IEEE 802.11-Based Ad Hoc Networks. In *Proc. of IEEE MASS*, pages 517–520, October 2006.

[131] H. Schwarz, D. Marpe, and T. Wiegand. Overview of the Scalable Video Coding Extension of the H.264/AVC Standard. In *IEEE Trans. Circuits Syst. Video Technol.*, pages 1103–1120, September 2007.

[132] S. Sharma. Analysis of 802.11b MAC: A QoS, Fairness, and Performance Perspective. Technical Report arXiv:cs/0411017v1, Department of Computer Science, Stony Brook University, November 2004.

[133] S. Shenker. Some conjectures on the behavior of acknowledgement-based transmission control of random access communication channels. In *Proc. of ACM SIGMETRICS*, pages 245–255, June 1987.

[134] M. Shreedhar and G. Varghese. Efficient Fair Queueing Using Deficit Round Robin. In *Proc. of ACM SIGCOMM*, pages 231–242, October 1995.

[135] A. Shriram and J. Kaur. A Measurement Study of Available Bandwidth Estimation Tools. In *Proc. of AMC IMC*, pages 39–41, October 2003.

[136] D. Stiliadis and A. Varma. Latency-Rate Servers: A General Model for Analysis of Traffic Scheduling Algorithms. *IEEE/ACM Trans. Networking*, 6(5):611–624, October 1998.

[137] J. Strauss, D. Katabi, and F. Kaashoek. Empirical Evaluation of Techniques for Measuring Available Bandwidth. In *Proc. of IEE INFOCOM*, pages 2162–2170, May 2007.

[138] W. Sun and K. G. Shin. End-to-End Delay Bounds for Traffic Aggregates Under Guaranteed-Rate Scheduling Algorithms. *Networking, IEEE/ACM Transactions on*, 13(5):1188–1201, October 2005.

[139] A. Tanenbaum. *Computer Networks*. Pearson Prentice Hall, 2002.

[140] F. A. Tobagi and L. Kleinrock. Packet Switching in Radio Channels: Part II–The Hidden Terminal Problem in Carrier Sense Multiple-Access and the Busy-Tone Solution. *IEEE Trans. Commun.*, 23(12): 1417–1433, December 1975.

[141] N. Vaidya, P. Bahl, and S. Gupta. Distributed Fair Scheduling in a Wireless LAN. In *Proc. of ACM MobiCom*, pages 167–178, September 2000.

[142] G. Verticale and P. Giacomazzi. An Analytical Expression for Service Curves of Fading Channels. In *Proc. of IEEE GlobeCom*, pages 1–6, December 2009.

[143] B. Walke, S. Mangold, and L. Berlemann. *IEEE 802.11 Wireless Systems - Protocols, Multi-hop Mesh, Relaying, Performance and Spectrum Coexistence*. John Wiley & Sons Inc., 2008.

[144] W. Wang, S. C. Liew, and V. O. K. Li. Solutions to Performance Problems in VoIP over a 802.11 Wireless LAN. *IEEE Transactions on Vehicular Technology*, 54(1):366–384, January 2005.

[145] D. Wu and R. Negi. Effective Capacity: A Wireless Link Model for Support of Quality of Service. *IEEE Trans. Wireless Communications*, 2 (4):630–643, July 2003.

[146] G. Xylomenos and G. C. Polyzos. TCP and UDP performance over a wireless LAN. In *Proc. of IEE INFOCOM*, pages 439–446, March 1999.

[147] G. Xylomenos, G. C. Polyzos, P. Mahonen, and M. Saaranen. TCP Performance Issues over Wireless Links. *IEEE Communications Magazine*, 39(4):52–58, April 2001.

[148] R. D. Yates and D. Goodman. *Probability and Stochastic Processes: A Friendly Introduction for Electrical and Computer Engineers*. John Wiley & Sons Inc., 2004.

[149] Y. Zhu, C. Dovrolis, and M. Ammar. Dynamic overlay routing based on available bandwidth estimation: A simulation study. *Computer Networks*, 50(6):742–762, April 2006.

OWN PUBLICATIONS

Zdravko Bozakov and Michael Bredel. SSHLauncher - A Tool for Experiment Automation in Distributed Environments. Technical Report KOM-TR-2008-11, KOM, TU-Darmstadt, July 2008.

Zdravko Bozakov and Michael Bredel. Online Estimation of Available Bandwidth and Fair Share Using Kalman Filtering. In *Proc. of Networking*, pages 548–561, May 2009.

Michael Bredel and Martin Bergner. On the Accuracy of IEEE 802.11g Wireless LAN Simulations using OMNeT++. In *Proc. of SimuTools*, pages 81–86, March 2009.

Michael Bredel and Markus Fidler. A Measurement Study of Bandwidth Estimation in IEEE 802.11g Wireless LANs using the DCF. In *Proc. of IFIP Networking*, pages 314–325, May 2008.

Michael Bredel and Markus Fidler. Understanding Fairness and its Impact on Quality of Service in IEEE 802.11. In *Proc. of IEEE Infocom*, pages 1098–1106, April 2009.

Michael Bredel and Markus Fidler. A Measurement Study regarding Quality of Service and its Impact on Multiplayer Online Games. In *Proc. of NetGames*, pages 1–6, November 2010.

Michael Bredel, Zdravko Bozakov, and Yuming Jiang. Analyzing Router Performance using Network Calculus with External Measurements. In *Proc. of IWQoS*, pages 1–9, June 2010.

Tronje Krop, Michael Bredel, Matthias Hollick, and Ralf Steinmetz. JiST/-MobNet: Combined Simulation, Emulation, and Real-World Testbed for Ad-Hoc Networks. In *Proc. of ACM WinTECH*, pages 27–34, September 2007.

SCIENTIFIC CAREER

Curriculum Vitae	
Name	Bredel, Michael
Day of birth	January 7th 1979
Education	
03/2009 – 02/2012	Research assistant and PhD student at the Institute of Communications Technology at Leibniz Universität Hannover.
	Research in the field of probabilistic performance bounds of communication networks. The project has been funded by the German research council (DFG).
05/2007 – 02/2009	Research assistant and PhD student at the Institute of Multimedia Communications at Technische Universität Darmstadt.
	Research in the field of probabilistic performance bounds of communication networks. The project has been funded by the German research council (DFG).
10/1999 – 04/2007	Studies of eletrical engineering and information technology at Technische Universität Darmstadt.

	Diploma thesis at the Institute of Multimedia Communications at Technische Universität Darmstadt, titled *"Design and Implementation of a Real-World Ad-Hoc Network Testbed based on JiST/MobNet"*
	Student project at the Institute of Microelectronic Systems at Technische Universität Darmstadt, titled *"Design und Implementierung einer automatischen VHDL-Codegenerierung aus optimierten Datenflussgraphen"*
07/1989 – 06/1998	Secondary school Alfred-Delp-Schule Hargesheim
07/1985 – 06/1989	Primary school Rümmelsheim

Working Experience

05/2006 – 08/2006	Student assistant at the Institute of Elektronic Systems, Technische Universität Darmstadt
01/2003 – 02/2006	Student assistant at the Institute of Microwave Engineering, Technische Universität Darmstadt
03/2002 – 12/2002	Student trainee at IP-Value in the field of software development and software tests
07/2000 – 02/2002	Student trainee at Siemens I&S IPS Frankfurt in the field of industrial purchasing services

Teaching Experience

since 04/2009	Hands-on exercises on networks and protocols
since 06/2007	Co-advising several final student projects and theses
04/2008 – 09/2008	Hands-on exercises on teletraffic theory