

Analysis and Design of Advanced
Multiservice Networks Supporting
Mobility, Multimedia, and
Internetworking
COST-279 – Midterm Report

COST-279 Management Committee
Michael Menth (eds.)

January 2004

Editors

COST-279 Management Committee

Chairman: José Brázio, Telecommunications Inst./IST, Lisbon, Portugal

Vice-Chairman: Phuoc Tran-Gia, University of Würzburg, Germany

Michael Menth, University of Würzburg, Germany

Chapter Editors and Contributors¹

Control Mechanisms for the Next Generation Networks

Editor: M. Menth

Contributors: H. Tarasiuk, S. Oueslati, P. Olivier, M. Menth, I. Gojmerac,
M. Meo, M. Caglar, Ö. Özkasap, I. Norros

Traffic Measurement and Characterization

Editor: M. Fiedler

Contributors: M. Caglar, P. Carlsson, M. Fiedler, G. Hu, J. Kilpi, U. Krieger,
P. Mannersalo, S. Molnár, L. Muscariello, P. Olivier, Ö. Özkasap,
A. Popescu, M.-A. Remiche, K. Salamatian, P. Salvador, K. Tutschku,
H. van den Berg, S. Wittevrongel

Queuing Theory

Editor: S. Wittevrongel

Contributors: N. Akar, C. Belo, M. Fiedler, D. Fiems, P. Gao, V. Inghel-
brecht, U. Krieger, K. Laevens, M. Mandjes, I. Norros, O. Østerbø, D. Sass,
A. da Silva Soares, K. Spaey, H. Tran, H. van den Berg, J. Virtamo, J. Wal-
raevens, S. Wittevrongel

Wireless Networks

Editors: W. Burakowski, A. Beben

Contributors: S. Aalto, vH. van den Berg, R. Boucherie, L. Cerdà, S. Imre,
J. Kilpi, M. Meo, M. Mohorčič, D. Staehle

COST-279 – Midterm Report

*Analysis and Design of Advanced Multiservice Networks Supporting Mobility,
Multimedia, and Internetworking*

Typeset in L^AT_EX

© Copyright January 2004

¹Please see chapters for detailed affiliations.

Preface

The present booklet is a report on the two first years of activity of COST Action 279, “Analysis and Design of Advanced Multiservice Networks supporting Mobility, Multimedia, and Internetworking,” prepared as a companion to its Mid-Term Seminar of January 21–22, 2004, in Rome, Italy. Following a short introduction to the framework, mode of operation, and activities of COST 279, the main part of the booklet gives a guided tour of the technical work developed within the Action, organized around the four main topics of *Control Mechanisms for Next Generation Networks*, *Traffic Measurement and Characterization*, *Queueing Models*, and *Wireless Networking*.

COST 279 is one of the Actions of the European COST Programme, an intergovernmental framework for European Co-operation in the field of Scientific and Technical Research allowing the co-ordination of nationally funded research on a European level. COST Actions are launched on a “bottom-up approach, with the initiative coming from the scientists and technical experts themselves and from those with a direct interest in furthering international collaboration. According to the general spirit of flexibility of COST, country participation follows an *à la carte* principle.

The COST framework provides means for the setting up of regular meetings among researchers of the participating countries, for the purpose of technical exchanges, discussion of research directions, and organization of common initiatives. The resulting co-operation and interaction among researchers is intended to help Europe hold a strong position in the field of scientific and technical research. Its experience has shown very beneficial to the research community: besides helping keep its cohesion, it has allowed advanced students and young researchers in the beginning of their careers to get in touch

with the latest research developments and integrated into the community, and it has potentiated the start up of research groups in environments where previous research tradition in an area does not exist. In addition, COST provides mechanisms for the transfer of its research results to the surrounding society, and examples abound of related “success stories.” Detailed information on the COST Program can be obtained via the Web at the URL <http://www.cordis.lu/cost/>.

COST 279 belongs to a distinguished lineage of COST Actions (COST 201, 214, 224, 242, and 257) that, since 1979, have fostered cooperation between European researchers in the field of teletraffic and multiservice network performance and design. This sequence of Actions has been carried on by what can be called a core group of the European teletraffic community, its members originating from industry, operators, and universities, and has been an ever-present actor in all the steps of the evolution of modern (digital) telecommunications networks—from circuit-switched networks to the introduction of ISDN and ATM in the 1980’s and on to today’s drive towards all-IP networks, from fixed to mobile and wireless service, from copper-based to fiber-optical transmission, and from electronic to optical switching. It thus has all the richness of the memory and experience of the past, together with the promise and potential for the future of its active members and permanently joining new researchers.

The [Memorandum of Understanding](#) (MoU) that established COST 279 states its objectives in the following way: “The main objective of the Action is to develop techniques for the analysis, design and control of advanced multiservice networks supporting mobility, multimedia and interworking, by means of the development and application of new and better analytical techniques for the mathematical understanding and optimisation of the behaviour of communications equipment, protocols, and network topologies and architectures, and of economic aspects such as pricing principles and network cost estimation. The results will have the form of mathematical models and results, algorithms, computer tools, and analyses of empirical traffic and network data. The Action is also expected to contribute to general, theoretical progress in basic issues like queueing theory, estimation and identification in stochastic models, and simulation, in particular as applied to rare events.”

For the achievement of its objectives the COST 279 community recognizes the importance of methodological progress regarding the analysis, design and planning of multiservice networks supporting advanced requirements. Therefore, a main distinguishing feature of the scientific programme of the Action

resides on the strong emphasis on modelling and performance issues, as pertaining to both the basic support disciplines and to their application to the study of the behavior of and interaction among networks, end-systems, services, and user applications.

COST 279 started on July 1, 2001, at a time where the research on ATM had been phased-out, and had been phased-in the research resulting from the widespread interest, arisen a few years before, of the use of the Internetwork Protocol (IP) for the operation and interconnection of heterogeneous networks. The Action has since steadily contributed to the understanding of these systems and provided inputs for their design, by means of work on topics such as techniques for traffic measurement, characterization, modelling, and estimation, network dimensioning, traffic control and engineering, and support of Quality-of-Service, all of these in the context of both wired and wireless networks and their interconnection. In parallel, the Action has continued to produce developments in basic issues like queuing models, statistical estimation, and random graph theory. The collection of these results is embodied in the 113 internal documents (officially know as *TDs*, shorthand for *Temporary Documents*) listed at the end of this booklet and available in the accompanying CD. From the work contained in these TDs about 110 publications have resulted in the peer-reviewed literature, the references of which are included in the Appendix *Bibliography*. The next four chapters of this booklet systematize the associated results and conclusions.

The core of the operation of COST 279 takes place at its Management Committee (MC) Meetings, occurring regularly at four-month intervals. In addition to the administrative and policy management aspects of the Action, the meetings include technical sessions where the internal technical documents are presented and discussed in an informal environment, with room for degrees of work maturation going from exploratory ideas, to work in progress, and all the way to finished work. The MC is formed by the following National Delegates of the 20 countries that have signed the MoU: T. Ziegler (A), H. Bruneel (B), G. Latouche (B), S. Louca (CY), A. Pitsillides (CY), U. Krieger (D), P. Tran-Gia (D), V. B. Iversen (DK), M. Villén-Altamirano (E), I. Norros (FIN), J. Virtamo (FIN), P. Olivier (F), K. Salamatian (F), P. Key (GB), S. Imre (H), S. Molnár (H), B. Zovko-Cihlar (HR), A. Baiocchi (I), M. Meo (I), T. Jensen (N), M. Rohne (N), J. van den Berg (NL), R. Litjens (NL), J. Brázio (P), R. Valadas (P), W. Burakowski (PL), G. Kandus (SI), M. Klimo (SK), P. Podhradsky

(SK), M. Fiedler (S), A. Popescu (S), and N. Akar (TR).

The meetings taken place so far have been attended by 135 different participants, coming from the following 39 research institutions from 18 countries: Telecommunications Research Center Vienna (A), Université Libre de Bruxelles (B), University of Antwerp (B), Ghent University (B), University of Cyprus (CY), University of Würzburg (D), Institute of Communication Networks, TU München (D), University of Stuttgart (D), Siemens, AG (D), Otto Friedrich University Bamberg (D), Technical University of Denmark (DK), Telefónica I&D (E), Universitat Politècnica de Catalunya (E), University of Paris VI—LIP 6 (F), France Télécom R&D (F), Helsinki University of Technology (FIN), VTT Information Technology (FIN), Budapest University of Technology and Economics (H), University of Zagreb (HR), University of Roma “La Sapienza” (I), Politecnico di Torino (I), Telenor Research and Development (N), TNO Telecom (NL), University of Twente (NL), KPN Research (NL), Delft University (NL), Centre for Mathematics and Computer Science (NL), Telecommunications Institute (P), Instituto Superior Técnico (P), Aveiro University (P), Warsaw University of Technology (PL), Blekinge Institute of Technology (S), Ericsson Core Network Development (S), Lund University (S), Telia Research (S), Institute Jozef Stefan (SI), University of Zilina (SL), Bilkent University (TR), and Koc University (TR). Further details about the operation of COST 279 can be obtained at the URL <http://www.lx.it.pt/cost279/>.

In addition to its “internal” operation, COST 279 implements activities to transfer its know-how to the outside community. One such activity is the organization of two Seminars, at the mid- and end-points of its operation, where its results are publicly presented. Another one is the implementation of a yearly Summer School specifically targeted for qualifying advanced students and young researchers of European research institutions in the field, with the stated objective of contributing to further strengthening the joint European research community. The first *COST 279 Summer School*, on *Stochastic Modeling and Analysis in Telecommunication* with emphasis on *Wireless Systems*, took place in August 2002, and was attended by 32 doctoral students from 20 research institutions belonging to 15 COST countries. The second edition of the Summer School took place in September 2003 having as topic *Routing and Multi-Layer Traffic Engineering in Next Generation IP Networks*, and was attended by 34 students coming from 27 research institutions of 18 COST countries. In both cases, lectures were provided by top European experts in the field,

pertaining to both universities and telecommunications operators. The initiative of a regular Summer School, not frequent among COST Actions, started as a learning experiment and has been quickly recognized as a success to be continued.

The achievements of COST 279 would not have been possible without the contribution of an enthusiastic and hard-working team of people. For the present booklet was instrumental the work of its editor, Michael Menth (University of Würzburg, Germany), who undertook the Herculean task of putting together all the needed typesetting infrastructure and of coordinating with the fine team of technical chapter editors the integration and edition of the four technical chapters. These were made possible by the indispensable and inestimable help of the contributing members listed at the beginning of each chapter. The Mid-Term Seminar is indebted to also the chapter editors, who organized the technical sessions, to Villy Bæk Iversen (Technical University of Denmark), who organized the *Optical Networking* session, to Phuoc Tran-Gia and Kurt Tutschku (University of Würzburg), who organized the Panel on *Peer-to-Peer Networking*, to Andrea Baiocchi (University of Rome “La Sapienza,” Italy), who provided the local organization, and to Kavé Salamatian (University of Paris VI, France), who provided key invited speakers contacts. The Summer School is a brainchild of Udo Krieger (Otto Friedrich University Bamberg, Germany), who single-handedly has been running all its logistical and technical aspects. The overall operation of COST 279 has had the wise guidance of its Management Committee, and its non-trivial administrative operation has been possible by the support of the External Scientific COST Secretariat and the work of the Actions Secretaries, Daniel Zaragoza and Ana Figueira (Telecommunications Institute—Lisbon, Portugal). And, last but not least, none of this would have been possible without the work of all the contributing members of COST 279, who are indeed its reason for existence.

January 2004

José Brázio
Chairperson, COST 279
Telecommunications Institute/IST
Lisbon
Portugal

Contents

1	Next Generation Networks	13
1.1	Introduction	13
1.2	Admission Control in Modern IP-Based QoS Networks	15
1.2.1	AC methods in the AQUILA IP-Based QoS Network	16
1.2.2	AC methods for Streaming Traffic	17
1.2.3	AC Methods for Elastic Traffic	20
1.2.4	Internet Service with Relative QoS	22
1.3	Network Admission Control and Resilience Aspects	23
1.3.1	Introduction to NAC Methods	23
1.3.2	The Performance of NAC Methods	24
1.3.3	Capacity Renegotiation	25
1.3.4	The Performance of NAC Methods under Resilience Requirements	26
1.3.5	Routing Optimization under Resilience Requirements	26
1.4	Capacity Dimensioning and TCP Dynamics	27
1.4.1	Access Network Links	28
1.4.2	Backbone Network Links	28
1.4.3	Streaming and Elastic Traffic Integration	32
1.4.4	Packet Level Issues	34
1.5	Routing Optimization and Load Balancing	36
1.5.1	Traffic Engineering Approaches Based on IP Routing	37
1.5.2	Traffic Engineering and Resilience Approaches Based on Multi-Protocol Label Switching	40
1.6	Scheduling	42

1.7	Multicast	44
1.8	Graph Models for Interconnection Networks	47
2	Traffic Measurement and Characterization	49
2.1	Introduction	49
2.2	Traffic Measurement	50
2.2.1	Passive Measurements	50
2.2.2	Active Measurements	52
2.3	Traffic Characterization	53
2.3.1	Multicast traffic	53
2.3.2	Optical Burst Switching traffic	55
2.3.3	TCP traffic	55
2.3.4	Modem pool traffic	56
2.3.5	GPRS traffic	57
2.3.6	WLAN traffic	57
2.3.7	P2P traffic	58
2.4	Traffic Modeling	61
2.4.1	PDF Estimation	61
2.4.2	Poisson- and Markov-type models	62
2.4.3	Time-discrete models	64
2.4.4	Flow-level models	65
2.4.5	Fluid models	67
2.4.6	Fractal-type models	69
2.5	Traffic Estimation	70
2.5.1	Traffic Inference	70
2.5.2	Traffic Prediction	72
3	Queueing Models	75
3.1	Introduction	75
3.2	Discrete-time queueing models	76
3.3	Fluid flow models	82
3.4	Gaussian storages	87
3.5	Processor sharing models	90
3.6	Other continuous-time queueing models	93
3.7	Queueing networks	96

4	Wireless Networks	99
4.1	Introduction	99
4.2	GSM Networks	100
4.2.1	Traffic Measurements and Modeling	101
4.2.2	Resource Management	102
4.3	UMTS Networks	104
4.3.1	Admission Control and Capacity of WCDMA Systems	105
4.3.2	Soft Handover	108
4.3.3	Impact of mobility on UMTS network performance	110
4.3.4	Scheduling of Packet Data Traffic in UMTS Networks	112
4.4	Wireless LANs	114
4.4.1	Performance analysis	115
4.4.2	Handover issues	115
4.4.3	Reliability aspects of mobile IP	119
4.5	Satellite communication	119
	List of Abbreviations	121
	Bibliography	127
	List of Temporary Documents	149

Chapter 1

Control Mechanisms for the Next Generation Networks

Michael Menth

University of Würzburg, Germany

Contributors:

Halina Tarasiuk (Warsaw University of Technology, Poland), Sara Oueslati (France Telecom, France), Michael Menth (University of Würzburg, Germany), Philippe Olivier (France Telecom, France), Ivan Gojmerac (FTW Vienna, Austria), Michela Meo (Politecnico di Torino, Italy), Mine Caglar, (Koc University, Turkey), Öznur Özkasap (Koc University, Turkey), Ilkka Norros (VTT, Finland)

1.1 Introduction

The Internet provides interconnection of networks based on the Internet Protocol (IP). In contrast to the structure of the telephone network—circuit-switched and connection-oriented—the Internet is packet-switched and connectionless. It is also unreliable, with end-to-end reliability provided to applications by the Transmission Control Protocol (TCP). TCP peers are located only at the end systems, leaving the core of the Internet stateless and the network unaware of higher protocol layers (e.g., service layer). Therefore, it is simpler to deploy a new service in the Internet than in the telephone network. The telephone world, however, has a revenue creating property not provided by IP networks: Quality of Service (QoS). This feature must be offered by a next generation of the Internet if circuit switched networks are to be integrated in or replaced by packet-switched networks.

The QoS feature can be implemented, e.g., by Admission Control (AC) or by bandwidth overprovisioning. AC requires AC entities in the network, raising interoperability issues and causing upgrade costs for systems. AC requires state storage in the networks, thus introducing the possibility for failure and potentially compromising the network scalability. These problems are avoided by overprovisioning but this approach has disadvantages: it leads to permanent increased bandwidth costs and the needed overdimensioning factor is not known.

Hence, the technical future of the Next Generation Network (NGN) is not yet decided and some of the research carried out in COST 279 contributes to the decision-making process in the sense that both directions are investigated.

Some of the work on the topic under COST 279 was done in two major research and development projects for Next Generation Networks.

- The first project is called AQUILA (Adaptive Resource Control for QoS using an IP-based Layered Architecture). It is a European research project partially funded by the Information Science and Technology (IST) Programme (number IST-1999-10077), with start date of January 1, 2000, and end date of March 30, 2003, for a total duration of 39 months. More details are given on the website:
<http://www-st.inf.tu-dresden.de/aquila/>.
- The second project is called KING (Key components for the Internet of the Next Generation). The overall objective of KING is to develop efficient solutions for carrier-grade IP networks that satisfy high QoS and resilience requirements by means of a common approach [1], while at the same time providing low operational overheads. Siemens Information and Communication Networks (ICN) has been running the KING research project since October 2001, together with a number of German universities and research institutes. The project is partially funded by the German Ministry for Education and Research (BMBF) for a period of 3 years. A website can be found at:
<http://www.siemens.com/king>.

Their results are partly presented in the first two sections of this chapter. Section 1.2 starts with practical applications of well-known results for AC and shows their feasibility. So far, AC has been investigated for unresponsive traffic, e.g., real-time traffic, whereas AC approaches for elastic TCP traffic can adapt to bottlenecks. Section 1.3 extends the concept of AC from link AC

(LAC) for a single link to network AC (NAC), which addresses the storage of AC information, i.e., the reservation states. Since failure recovery is harder for stateful networks, resilience issues are studied in this context.

Section 1.4 is dedicated to the dimensioning of networks. Today's IP networks have to be provisioned reasonably to get satisfying QoS values for elastic traffic, e.g., response time for TCP-based applications such as web browsing. Then, the amount of required capacity is investigated with regard to real-time QoS measures like loss and delay.

The remaining part of this chapter is dedicated to network improvements. Routing optimization and load balancing (cf. Section 1.5) avoids overload and thereby improves QoS or, conversely, allows network operation at a given QoS level with less resources. Similar benefits are obtained with suitable scheduling mechanisms (cf. Section 1.6) because they can decrease the queuing time for delay sensitive data. Concurrently, multicast (cf. Section 1.7) contributes to network efficiency because of shared transmissions reducing the traffic volume.

The results obtained for most of the issues mentioned depend on the network topology. Therefore, Section 1.8 is dedicated to network graph models.

1.2 Admission Control in Modern IP-Based QoS Networks

The modern IP-based QoS networks are targeted for effective handling of traffic produced by a variety of application types currently available to the Internet users. These applications differ in transmission requirements, demanding from the network the support of a number of network services differing in QoS objectives. The design process of a network service assumes that a specific traffic profile should be handled by the network with a predefined QoS, usually expressed by such parameters as packet transfer delay and packet loss ratio. To meet these requirements, some QoS mechanisms in different levels of the network should be implemented. Mechanisms like classifiers, conditioners, and schedulers are used at the packet level, while AC is performed at the flow level.

We classify the QoS architectures with regard to strict or relative QoS. For providing strict QoS, the AC function constitutes a key element, allowing the regulation of the volume of traffic submitted to a network with limited

resources. On the contrary, the relative (proportional) QoS does not use AC but rather assumes that high priority flows are handled by the network in a preferred way. This section contains a brief overview of network service types that are recognized as adequate for networks with strict and relative QoS.

In the last decade many AC methods have been developed for traffic control in Asynchronous Transfer Mode (ATM) networks. Both ATM and IP networks are packet-switched technologies with many similarities. Therefore, it seems natural to adopt the AC methods for ATM to IP technology but not without considering their differences. Some important issues are:

- The fixed size of packets (cells) in ATM has the need for a lower level of preventive AC compared to IP where packets have different size. As a consequence, more packet jitter and delay is expected for IP technology. This leads to serious problems for efficient multiplexing of traffic with different profiles. For instance, small voice packets of 60 bytes experience relatively large delay when they wait for the end of transmission of a 1500 bytes data packet.
- Effective AC mechanisms were developed for streaming traffic in ATM. These results are not applicable to elastic TCP-controlled traffic, such that new AC methods are required for IP networks.

AC methods can be classified into two main types: Traffic Descriptor Based AC (TDAC) methods and Measurement Based AC (MBAC) methods. TDAC methods can be again split into deterministic and statistical multiplexing schemes. Furthermore, both statistical and MBAC methods can refer to Rate Envelope Multiplexing (REM) or to Rate Sharing Multiplexing (RSM). The above classification, earlier introduced for ATM, can be extended to IP.

1.2.1 AC methods in the AQUILA IP-Based QoS Network

A framework for providing QoS differentiation and strict QoS guarantees in IP networks is presented in [2]. This approach assumes network support for four premium network services and is based on the Differentiated Services concept. Each network service is designed to support a class of applications with similar QoS requirements and traffic characteristics. For each premium network service an adequate TDAC method is proposed based on AC algorithms developed for ATM.

The AC for the Premium Constant Bit Rate (PCBR) network service is based on the REM scheme using a well-known peak rate allocation method. The admissible load for the capacity allocated to PCBR is calculated based on the analysis of an M/D/1/B system depending on the assumed target packet loss ratio and the available buffer size. PCBR is dedicated to constant bit rate traffic (e.g., voice trunks) and is served with the highest priority. Negligible packet delay variation can be assured when at most 10% PCBR load is allowed on the link.

The REM scheme is also applied to the Premium Variable Bit Rate (PVBR) network service, which is designed for effective transfer of streaming variable bit rate traffic, e.g., video applications. Here, the effective bandwidth is calculated by Karl Lindberger's method [3]. The buffer dimensioning rules are adequate for absorbing the so-called packet scale congestion. For this purpose, the analysis of the $N \cdot D/D/1$ queuing system was applied.

The Premium Multi Media (PMM) network service is used for greedy TCP (or TCP like) flows. Its AC method is a function of a sustained bit rate value which corresponds to the minimum required flow rate.

Finally, the Premium Mission Critical (PMC) network service is conceived to assure few losses for elastic non-greedy flows. Its AC method adapts the RSM scheme and the effective bandwidth calculation.

In addition, the Standard network service is designed for carrying best effort traffic.

An extension of MBAC methods for streaming flows and two novel AC concepts for PMM service are provided in [4, 5]. The final results of the AQUILA project confirm the expected efficiency of the proposed AC methods for the proposed premium network services. They show the feasibility of differentiated QoS network services by adding new functionality into existing IP networks. Selected results are presented, e.g., in [6].

1.2.2 AC methods for Streaming Traffic

In the following we present a two-stage AC method for voice traffic, an MBAC approach for non-real-time variable bitrate traffic, and an application of AC results to investigate protocol design alternatives for low bitrate real-time traffic. These studies are based on well-known AC methods for ATM networks. They

extend these methods by taking the properties of IP networks into account.

A Two-Stage AC Method

As mentioned before, IP networks serve packets with different sizes. This nature of IP traffic introduces additional problems for delay and jitter guarantees for streaming (real-time) traffic compared to ATM networks. A possible solution has been proposed in [7]. A two-stage Connection Admission Control (CAC) model based on REM and Simplified Reference Model (SiRM) concepts is proposed and investigated for voice traffic over IP networks. The CAC method is able to deal with heterogeneous voice traffic flows of ON-OFF nature with loss and delay jitter requirements. AC decisions in each module are made as follows. If a new voice flow does not make the rate overloading probability increase beyond a predefined threshold, it is admitted by the first module and passed to the second module. The AC decision of the second module is based on criteria related to delay and loss metrics. The loss metric is the buffer overflow probability. For the delay metric, either the maximum queuing delay, or the mean queuing delay, or a given percentile of the queuing delay is chosen as the QoS criterion. The delay percentile should be adopted as a criterion because it is more meaningful in the context of statistical end-to-end jitter guarantees. One concrete realization of the proposed approach is studied, where the REM scheme is ensured by the well-known Chernoff bound based effective bandwidth concept. The model of SiRM is obtained by exploiting the recently proposed Negligible Jitter (NJ) conjecture [8]. The application of the NJ conjecture coupled with the non-preemptive scheduling scheme (with high priority service for voice traffic and low priority for best effort traffic) leads to the implication of an M/D/1/K queue with exhaustive service and multiple vacations. The vacations of the server correspond to a situation where the output link is occupied by the best effort traffic. Analytical results show that loss and delay metrics computed with the theoretical model are quite close to those obtained with simulation, in particular when the mean packet size is used in the reference SiRM model, rather than the Maximum Transfer Unit (MTU) packet size suggested earlier by the original NJ conjecture. A potential extension of the proposed CAC scheme to a network is the application of the scheme in each network node and ReSerVation Protocol-like (RSVP) signaling for a network-wide CAC decision.

Measurement Based AC

In general, TDAC schemes aim at guaranteeing a QoS level for the traffic submitted in accordance with the declared profile. Their main weakness is the dependence of their efficiency on the suitability of the a-priori declared traffic parameters. However, since tight values are not known a priori, these parameters are usually overdimensioned. Therefore, the authors of [9] propose a new method for more precise evaluation of the traffic descriptors. The recognized approach for traffic characterization is based on the token bucket mechanism, which is described by two parameters: token accumulating rate and bucket size. Note that these parameters constitute a base for effective bandwidth evaluation and they determine the amount of resources the network should dedicate for handling the corresponding connection. The proposed method is based on online traffic monitoring, which allows to obtain the minimum effective bandwidth and to reduce the required resources. The approach is investigated for handling non-real-time variable bit rate traffic with zero packet loss as QoS requirement. In this case the effective bandwidth is calculated according to the Elwalid-Mitra-Wentworth (EMW) formula [10]. The proposed algorithm for assessment of the token bucket parameters is based on simultaneous traffic monitoring by a number of modified leaky bucket mechanisms. The effective bandwidth calculation takes into account the optimal point from the burstiness curve, which minimizes the EMW formula. The effectiveness of the proposed approach is illustrated by the simulations for different traffic traces of Moving Pictures Expert Group (MPEG) video and Local Area Network (LAN) traffic.

Application of AC Results for Performance Evaluation of Protocol Design Alternatives

AC results are also used for the investigation of protocol design alternatives [11]. Low-bitrate real-time traffic usually comes in small packets leading to an inefficient ratio of RTP/UDP/IP header and payload size. Such traffic is produced by voice, video, and circuit switched applications, e.g., in the terrestrial radio access networks of wireless communication systems like Global System for Mobile (GSM) or Universal Mobile Telecommunications System (UMTS). Base stations are connected to radio network controllers over low-bandwidth links because they produce usually only little traffic. As the fixed network capacity is expensive on leased lines, too, header compression techniques are

attractive to save bandwidth. The assumed QoS requirements for low-bitrate real-time data traffic are little packet loss and delay. The notion of the Delay Budget (DB) is introduced and the QoS requirement is that the probability of the packet waiting time exceeding the DB be smaller than a target value. A suitable AC mechanism for that scenario is based on the $N \cdot D/D/1$ queuing system, from which the packet waiting time distribution can be derived. This model requires homogeneous flows, i.e., flows with the same constant packet inter-arrival times and the same constant packet sizes. In addition, reasonably designed buffers can prevent packet loss. The analytical results in [11] compare the critical net load for voice traffic with and without header compression. The impact of header compression on the critical load is two-fold: both the mean bit-rate and the burstiness of the flows are reduced. As a consequence, 150% more traffic can be carried over the same low-bandwidth links.

1.2.3 AC Methods for Elastic Traffic

The AC methods for elastic traffic may be based on TDAC and MBAC. The first two algorithms hereafter described fall in the first category. They use the notion of traffic contract between user and network and require explicit resource reservation based on declared traffic descriptors, as opposed to the third approach which proposes an implicit MBAC algorithm.

The AC algorithm proposed in [5] was developed in the framework of the IST European project AQUILA. The algorithm is designed to fit within a Diff-serv architecture and to meet the QoS requirements of a network service that consists of long-lived greedy TCP connections such as those generated by File Transfer Protocol (FTP) users. Prior to establishing a TCP connection, the application sends a request to the AC agent which is located at the edge router, specifying a target rate. This approach relies on the notion of traffic contract that needs to be enforced by the edge routers. Each admitted TCP connection is subject to policing based on a token bucket mechanism whose parameters are inferred from the target rate and Round Trip Time (RTT) statistics (minimum and average if available). The token bucket parameters (token filling rate and bucket size) constitute input parameters for admission decision. The role of the AC agent is to make sure that the traffic offered to the core network does not exceed the dedicated resources (bandwidth and buffer) on the corresponding edge router to core router link. Precisely, a new TCP flow is accepted if the following conditions are satisfied: the sum of the token bucket rates (re-

spectively, bucket sizes) assigned to the ongoing flows and to the incoming flow must be less than or equal to the amount of capacity (respectively, buffer) dedicated to the corresponding service on the link. Details on this approach and numerical evaluation based on simulations can be found in [5].

Another AC algorithm [12] targets sporadic short-lived flows requiring a new service, so-called Premium Message Handling Service (MHS), in the context of a Diffserv network. This service is designed for successfully delivering short messages to destinations, using point-to-point TCP connections. The message delivery delay (transfer time) should not exceed a predefined threshold value. Handling such traffic assumes that messages have a fixed limited size. Based on this assumption, the authors compute the maximum number of packets that can be sent in one burst. The transmission of a whole message typically completes during the TCP slow start phase. At packet level, a separate queue for Premium MHS packets is required in each router with a dedicated amount of bandwidth and buffering. The buffer size should be dimensioned so as to avoid packet losses and satisfy the message transfer time requirement. At flow level, the admission decision is based on the maximum number of admitted connections. The corresponding threshold is inferred from the traffic descriptors declared by each TCP connection, the dedicated buffer size and bandwidth, and allowed message transfer time. Traffic descriptors consist of the peak bit rate, which in this case denotes the link rate connecting the host to the edge router, and the maximum burst size. The impact of message sizes on the maximum number of admitted flows is also investigated; the bigger the message size, the smaller the acceptance area.

In contrast to the two previous approaches, a third AC mechanism [13] applies to elastic flows that do not have a clearly defined service requirement. Here, the objective of AC is to avoid the negative situation where demand overload in the absence of AC by an increasing number of flows in progress results in ever smaller throughputs for each one of them. Demand overload can occur in any reasonably well dimensioned network for a variety of reasons, including link and/or router failures. The AC mechanism is implicit in that flows are not preceded by any explicit signalling exchange between user and network. The user generates flows spontaneously as in the current best effort architecture, and the network elements performing AC detect them “on the fly” by maintaining a list of *protected flows* and systematically comparing the flow identity of all arriving packets with this list. The flow identity is determined from in-

variant fields in the IP and TCP packet headers. In IPv6, the “flow label” field could be used in conjunction with the origin and destination IP addresses. In the absence of congestion, a new flow is added to the list of protected flows and the packet is forwarded. If the link is already too heavily loaded, the first packet (or packets) of a new flow is simply discarded. The current state of congestion is determined based on measurements since there are no declared flow traffic specifications. A new flow is accepted if the estimated available bandwidth is greater than or equal to a predetermined admission threshold. An optimal choice of admission threshold should produce negligible blocking under normal load while maintaining sufficiently high throughput for admitted flows in overload [14]. The authors underline that the solution sought does not need to be highly accurate; elastic traffic is particularly tolerant to estimation errors and MBAC is inherently self-correcting. To demonstrate the feasibility and efficiency of the proposed AC mechanism, a testbed was setup and trials were carried out with real traffic. The AC experiments in [13] were performed on a 10 Mbit/s Ethernet segment. New tools were developed to visualize realized flow level performance (bubble diagrams, box diagrams). They clearly illustrate the phase change in perceived performance occurring as demand evolves from normal load to overload. Preliminary test results of an upgraded system on a higher speed interface (a Gbit/s Ethernet) suggest that the solution is scalable and could be implemented without much difficulty in modern backbone routers.

1.2.4 Internet Service with Relative QoS

A solution for receiving relative QoS using DiffServ mechanisms is provided in [15]. The authors choose the Dynamic Real-Time/Non-Real-Time (RT/NRT) mechanism, which have received less attention than Assured Forwarding (AF) or Expedited Forwarding (EF), and related Simple Integrated Media Access (SIMA) proposal [16]. SIMA relies both on financial and end-to-end congestion control incentives to achieve differentiated bandwidth allocation. The entity that influences both charging and division of bandwidth is the nominal bitrate. AC is not considered. First, flow and packet level models for SIMA are presented and compared for elastic flows. Then, fair differentiation among elastic flows using SIMA is shown as well as the fair differentiation among TCP and non-TCP flows. Finally, the applicability of SIMA as a future Internet Service model is demonstrated. The obtained results illustrate that SIMA

does not result in fixed weights, determined by the nominal bitrate or price paid, between bandwidth allocations of different classes. However, the provided models showed that it did achieve differentiation between classes, where the weights varied between equal allocation to allocation in proportion to the nominal bitrate purchase.

1.3 Network Admission Control and Resilience Aspects

Traditionally, AC has been primarily understood to be the answer to the question: How much traffic can be carried over a single link without violating the QoS requirements in terms of packet loss and delay on that link? In the following we designate by link AC (LAC) this kind of AC. Examples for LAC are peak rate allocation, the concept of effective bandwidths and, to some extent, measurement based Admission Control. For practical implementation, the traffic limitation on all network links is required. This can be done in a straightforward way by applying LAC methods on a link-by-link basis, e.g., by using resource reservation protocols like RSVP. Such approach, however, induces information states about individual reservations inside the network if they are collocated with the routers. Another option is a bandwidth broker solution [17] that keeps all information in a centralized database. Both approaches are problematic. The first one requires core routers to be aware of AC decisions and the second one presents a single point of failure. Therefore, a NAC is desired that resides only at the border routers.

1.3.1 Introduction to NAC Methods

In [18] four basically different NAC approaches are given. Flows are admitted or rejected based on the capacity of virtual budgets they have to ask for admission. The size of these budgets must ensure that no congestion can occur on any link in the network [19].

- The link budget (LB) based NAC defines a capacity budget for each link and flows have to be admitted by all LBs that are associated with their paths. This is depicted in Figure 1.1(a) and corresponds to the

above described link-by-link application of LAC methods, which induces reservation states in the network or single points of failure. In contrast, the following NAC methods require reservation states only at the border routers.

- The ingress and egress budget (IB/EB) based NAC performs an AC decision only at the ingress and the egress router of a flow—independently of each other (cf. Figure 1.1(b)). This concept is known from the Differentiated Services context [20] and is implemented in the AQUILA project [2, 21].
- The border-to-border (b2b) budget (BBB) based NAC has a BBB for each b2b relationship, which may be consulted at the ingress border router as illustrated in Figure 1.1(c). In contrast to the IB/EB NAC, the choice of the BBB depends both on the ingress and the egress router of the corresponding flow request.
- Finally, there is the ingress and egress link budget (ILB/ELB) based NAC, where every ingress and egress border router holds a link specific budget for each link in the network and the flow asks for admission both at the ingress and egress router, consulting any budget related to the link on its path (cf. Figure 1.1(d)). If only ILBs are used, the NAC may be viewed as a local bandwidth broker disposing of a private share of the total network capacity and is similar to the hose model [22].

1.3.2 The Performance of NAC Methods

The average resource utilization in a suitably dimensioned network is a reasonable performance measure for NAC methods [23]. The LB NAC reveals the best performance, followed by the ILB/ELB NAC, and the BBB NAC. For little offered load in terms of average number of simultaneous flows the differences are significant, while for sufficiently large offered load the LB NAC, ILB/ELB NAC, and BBB NAC converge to 100% potential resource efficiency. In contrast, the IB/EB NAC has a significantly lower performance, since the NAC can not avoid pathological traffic patterns that cause congestion on some links and leave other links unused. In general, the performance depends on many aspects like the network topology [24] or the traffic matrix and the routing [25].




Figure 1.1: Budget based network admission control (NAC) methods.

1.3.3 Capacity Renegotiation

In case of the LB NAC the LBs can be used by any flow in the network traversing the respective link, while the other budget types can be used only by a subset of flows, e.g., by those entering the network at a common ingress router. Flow blocking occurs if the capacity of a budget for such a subset of flows is exhausted. This can happen although the network is still far from being overloaded, e.g., if other budgets are hardly used. In such case the network capacity is badly distributed to the budgets and the capacity should be renegotiated. The AQUILA architecture is conceived to perform that renegotiation in a layered and scalable manner [26] and algorithms to achieve the redistribution

are investigated [27]. Capacity renegotiation among the border routers costs some signalling efforts and should not be triggered for smallest capacity units. Hence, capacity budgets are increased by a chunk of bandwidth which leads to overreservation for traffic aggregates. This tradeoff between overreservation and signalling reduction is investigated in [28] and analytical results are presented.

1.3.4 The Performance of NAC Methods under Resilience Requirements

An important feature of a Next Generation Internet is the provisioning of reliable QoS services, i.e., the QoS should not be compromised in case of short-term local network failures [1]. If a local outage occurs, the traffic is quickly rerouted around the failed network element and is still delivered to its destination. The QoS can only be maintained if sufficient capacity is available on the deviation route. In this case the network is resilient to failures. To that aim, a set of potential failure patterns (e.g., all single link failures) together with the corresponding rerouting has to be taken into account when the capacities of the NAC budgets are set. The performance of the NAC methods is investigated under these side conditions in [29]. Under resilience requirements the BBB NAC is most efficient in terms of resource efficiency, followed by the ILB/ELB NAC and the LB NAC. The IB/EB NAC is still the least efficient. The bandwidth utilization also increases for all NAC methods with increasing offered load, but remains clearly below 70% even in the best case. For very large offered load the reciprocal of that value is the amount of additional required backup capacity, and it can be in the range from 50% to more than 100%. As the BBB NAC shows the best performance under resilience requirements, it is also implemented in the KING testbed.

1.3.5 Routing Optimization under Resilience Requirements

The results in [29] also show that the routing has a definite impact on the required backup capacity, in particular in networks with a high offered traffic load. A routing optimization problem under resilience requirements exists if backup resources can be shared by various flow aggregates in different failure

scenarios. In this case the required backup capacity is a performance measure that can be minimized. In the case of the BBB NAC, the impact of the NAC vanishes for large offered load and the routing optimization problem can be solved for static aggregate sizes as indicated in the traffic matrix. In [30] the concept of path protection and self-protecting multi-paths are suggested. They are both b2b protection mechanisms and can be implemented by label switched paths (LSPs) in MPLS. An optimization algorithm for the layout of these structures and for the load balancing of the multi-paths is presented. The results show that only 17% backup capacity is required in some networks to protect them against all single link and router failures using the self-protecting multi-paths approach. The amount of required backup capacity depends mostly on the network topology. Strongly meshed networks favor the existence of link- and node-disjoint multi-paths and contribute to cost-efficient solutions.

1.4 Capacity Dimensioning and TCP Dynamics

To correctly dimension network links it is essential to understand the three-way relation between link capacity, expressed demand, and realized quality of service. Given the complexity of Internet traffic at the packet level (self-similarity, long range dependence, etc.), such a relation is very difficult to obtain in this domain. So it is very appealing to develop performance models of elastic (data) traffic at a higher level, namely the traffic “flows” or “sessions”. One can refer to [31], for example, to get a thorough discussion on the concept of a flow and some possible definitions. In this case it is appropriate to express dimensioning objectives in terms of the mean flow throughput or equivalently the mean transfer time as perceived by users to download some digital documents. Assuming a Poisson flow arrival process (for a source population of infinite size) and fair sharing of the available bandwidth by concurrent flows (supposed to be an ideal achievement of TCP dynamic control), the performance models generally fall in the category of Processor Sharing (PS) queueing systems, which under very general assumptions exhibit the very nice and useful property of insensitivity [32, 33].

1.4.1 Access Network Links

For access networks, a typical example of the above cited three-way relation is the Engset loss formula for classical circuit-switched telephony. In [34] it is shown that similar relations exist for high speed IP access networks carrying data traffic. Assuming a source population of finite size and a fair sharing of the capacity among the user-generated concurrent flows, some “PS-generalized Engset” formulas are derived, relating capacity, demand expressed in terms of the per-user offered traffic, and performance quantified by the useful per-flow throughput. Performance is shown to be largely independent of precise traffic characteristics such as the statistical distributions of flow size and “think time” [35].

Very simple approximations of this exact model can be obtained, exhibiting two distinct performance regimes (leading to two corresponding strategies for dimensioning): a transparent regime where the mean useful rate equals the access rate; and a saturated regime where the whole capacity is always used and fairly shared by active users. Analysis of the impact of a possible heterogeneous demand demonstrates that a homogeneous demand constitutes a worst case, so that considering an average offered traffic would be a conservative strategy for dimensioning. From these results, the definition of simple dimensioning rules follows and is illustrated in Figure 1.2. Note that operating in the saturated regime, in order to meet a target useful rate lower than the access rate, might be quite risky since the performance is highly sensitive to the accuracy of the offered traffic estimate.

Some extensions to the approximate model are considered in [34], such as situations with unfair bandwidth sharing or different access rates. Broadly speaking, it remains that the key parameter for dimensioning is the offered traffic, defined as the average data rate a user would generate in the absence of congestion.

1.4.2 Backbone Network Links

Best Effort Provisioning

For dimensioning IP backbone networks, one possible approach is bandwidth (over-)provisioning, where network transparency is ensured without implementing any specific QoS mechanism. Here “transparency” means that traffic




Figure 1.2: Dimensioning function based on approximations - Access rate = 500 Kbit/s

flows are not constrained by the considered network links and can make full use of their access rates. In line with this approach, [36] suggests to make use of standard traffic measurements performed by network operators on a relatively large time scale (e. g. 5 or 15 minutes), while the time scale at which the QoS is perceived in terms of transparency is much finer (about 1 s). In this framework, the adequate QoS indicator is defined as the probability (fraction of time) that the aggregated rate of offered traffic is greater than the link rate.

Two main assumptions seem appropriate for backbone links with high level of traffic aggregation and high capacity with respect to common access rates: (i) the aggregated traffic intensity is Gaussian; (ii) the number of concurrent flows results from a fluid $M/G/\infty$ model [37]. Under these assumptions, theoretical expressions of the required capacity are derived in [36] as a function of the average (offered) load and the transparency time scale. A simple explicit formula can be obtained in case of exponential flow size distribution. An experimental validation of the method is performed in an operational network environment with Asymmetric Digital Subscriber Line (ADSL) traffic: as can be seen in Figure 1.3, the model is shown to provide a very good fit to the experimental data relating the long term average rate to the “peak rate”

expressed as the 99% quantile.




Figure 1.3: 99% quantile of 1 s traffic rate as a function of 15 min mean traffic rate

QoS Differentiation

In the perspective of dimensioning TCP/IP networks supporting QoS differentiation, the performance of Internet links carrying TCP traffic with two priority classes is considered in [38], see also [39]. The elastic flows are assumed to appear according to a Poisson process and to fairly share the capacity (link bandwidth). Hence the stochastic process of TCP flows is described by a

multi-server PS queueing model, $M/G/C$ -PS, with two customer classes. The “number of servers” C of the queueing system is the ratio of the link capacity over the access rate, assumed identical for all users.

Considering that high priority flows have strict priority (pre-emptive resume) over low priority flows, closed-form expressions are derived for the mean transfer time of high-priority flows in the general case, and of low-priority flows in some special cases with exponential flow size distribution [40]. In the general case, only approximate expressions for the mean transfer time of low-priority flows can be provided. By comparing with *ns* simulations, it is demonstrated that the model provides highly accurate performance predictions when the mean flow size is at least 10 IP-packets, the loss rate is negligible, and the total traffic load does not exceed 70-80% (see Figure 1.4). When those conditions are not met, some packet level TCP features such as slow-start and packet loss effects create significant discrepancies between simulation results and the simplified flow model outputs.




Figure 1.4: TCP performance for low priority customers - Pareto file size distributions; Link rate = 10 Mbit/s

1.4.3 Streaming and Elastic Traffic Integration

Although TCP is still the dominant transport protocol in the Internet (about 90% of the overall traffic, measured in bytes, according to recent experiments), it is not possible to ignore the current quick development of real-time applications (voice, video). So, in the near future, streaming and elastic traffic will co-exist in the Internet with non-negligible part for each of them, and the problem of their integration within shared resources has to be tackled. Two contributions in COST 279 have provided substantial advances in that direction, by investigating the performance of TCP data transfers in presence of a time-varying bandwidth. This variable capacity is that which is left available by streaming flows, the packets of which are supposed to have priority with respect to those of elastic flows.

Flow Level Dynamics

First, [41] concentrates on flow level TCP performance by considering a dynamic system where elastic flows as well as streaming flows appear and disappear according to stochastic processes. Assuming Poisson flow arrivals on a network link and fair bandwidth sharing among elastic flows, the model is based on the performance of an M/G/1 Processor Sharing queue with time-varying capacity. It is shown that the mean transfer time (or, equivalently, mean throughput) of elastic flows may be greatly unstable if the streaming traffic volume is not controlled. On the contrary, when the conditions for uniform stability are met, e.g., by performing admission control on the streaming calls, the available bandwidth is always greater than the elastic traffic demand so that realized throughput is good and more easily predictable.

In the case of uniform stability, stochastic bounds of the mean transfer time are given which are shown to be rather tight and to depend only on the mean streaming and elastic loads, being insensitive to other detailed traffic characteristics. These bounds correspond to two extreme regimes defined by the time scales at which the streaming flow and the elastic flow processes evolve, respectively. The fluid regime is a best case providing the lower bound for response time, where the streaming flow process evolves so fast that elastic traffic sees a constant available bandwidth. At the opposite, the quasi-stationary regime is a worst case providing the upper bound, where elastic flows see a succession of equilibrium states with constant number of streaming flows. A

set of simulation experiments illustrates how the elastic traffic performance evolves between these two bounds, as a function of elastic load and under different streaming or elastic flow size distributions. For example, Figure 1.5 shows, rather surprisingly, that in the local instability region the elastic flow throughput is higher when the variability of flow size is high; besides, system insensitivity is verified in the uniform stability region.




Figure 1.5: Impact of the elastic flow size distribution on the performance (elastic flow throughput) of streaming/elastic traffic integration

TCP Behavior

As a complementary approach, [42] deals with packet level dynamics of TCP connections: an analytical model is developed which is able to reproduce the main features of TCP (Reno) behavior in presence of time-varying available capacity. The interaction between one persistent TCP flow and high-priority real-time traffic (e.g., voice and video) is modeled by a two-dimensional continuous time Markov chain where the state vector is defined by the number of

current streaming flows and the size of the TCP congestion window.

The accuracy of the proposed model, as compared to extensive ns simulations, is shown to be satisfactory on the whole and especially good when the buffer size is of the same order as the connection bandwidth-delay product and when the time scale of streaming flow dynamics is much larger (one order of magnitude) than the RTT. Besides, as shown in Figure 1.6, the model provides better results than other popular TCP models which appear either optimistic [43] or pessimistic [44] in predicting the average flow throughput. Finally, among other interesting TCP features derived from the model and simulations, an important result (confirming earlier simulation works) is the relative inability of TCP to make full use of the available bandwidth when sharing the capacity with higher-priority traffic: an efficiency coefficient of about 70-80% is typically observed with respect to the ideal case of constant capacity and correctly sized buffer.

1.4.4 Packet Level Issues

Although not initially designed to do so, the family of TCP/IP protocols is widely used today, and with little doubt in the near future, to provide support for ultra high-speed communication technologies over a multimedia Internet network, covering the fields of computer communications as well as traditional communications. To complement the flow level approach for performance analysis, as assumed in previous sections, one may believe that proper adaptation of TCP/IP to this new telecommunications world and its requirements should be necessary and could be achieved in combination with Quality of Service and connection control mechanisms (DiffServ, RSVP, ...). To this end, [45] describes the issues and requirements regarding the TCP/IP protocol operation, particularly focusing on the effect of network buffering and on the analysis of Active Queue Management (AQM) schemes. Some simulations illustrate how AQM, specifically Random Early Detection (RED) and Explicit Congestion Notification (ECN), allow the avoidance of queue overflow and help managing the realized bandwidth and packet discard process in case of various flows with different priorities.

RED is one of the most popular AQM algorithms, widely studied in the Internet community, whose aim is to “stabilize” TCP dynamics in case of overload in the network. Such mechanism implemented in router buffers begins to drop packets in the early stages of congestion, according to a drop probability




Figure 1.6: TCP connection performance (throughput) from models and simulations - Streaming traffic (supposed VoIP) offered load = 0.6

function of the average queue size. It is an interesting challenge to discuss the proper shape the RED drop function should have. Based on simulation experiments, [46] shows that the standard linear drop function, considered in the original RED algorithm [47], cannot cope with a broad range of load. Particularly, significant under-utilization of link capacity is observed at low loads, due to the drop function linearity.

Given the mathematical constraints that are derived to meet the low load and high load requirements, [46] considers two ways for determining the proper shape of the drop function, based on a class of polynomial functions and on a TCP performance model. Performed simulations of FTP permanent connec-

tions and Web-like transfers demonstrate that certain properly derived drop functions allow a broader range of loads and yield less varying averaged queue size as compared to the original linear function.

Finally, in order to support Service Level Agreements on IP networks, some guidelines about resource dimensioning are given in [48]. Based on OP-NET simulations of some Internet services such as Web (HTTP) and e-mail (SMTP), the impact of the link layer (assumed Frame Relay) and transport (TCP) control parameters on packet delays and application response times is analyzed in details. Particularly, an interesting point is that the end-to-end packet delay along the Frame Relay Permanent Virtual Circuit (PVC) path mostly results from the queueing delay at the Wide Area Network (WAN) ingress points, due to the significant bandwidth differences between LAN and WAN link layers.

1.5 Routing Optimization and Load Balancing

Since the early days of IP networks, researchers have been working on problems related to the routing of traffic. The first major challenge was keeping the routes consistent and loop-free in the entire network domain even in the presence of equipment failures. In order to assure this, the focus of research and engineering was set on the automatic distribution of topology information throughout the network and on the automatic generation of routing tables. The result of these efforts was the intra-domain routing architecture of the Internet, which has remained practically unchanged until today. In this architecture, link weights (sometimes also called link costs) are assigned to each directed link in the domain and used for the calculation of paths. The paths between any two nodes in the domain are determined such that the sum of link weights (i.e., link costs) over all path candidates is minimized. These paths are called “minimum cost paths” or “shortest paths.” On top of this architecture, routing protocols have been developed, which allow automatic dissemination of topology information throughout the network. These protocols have been designed for achieving network robustness and fast routing re-convergence in the case of network faults. With this routing architecture the traffic paths remain static as long as the link weights are not changed, meaning that the routing of traffic

is not sensitive to load conditions in the network. As the goal of robustness has been achieved, the focus of research has recently shifted to traffic engineering, i.e., performance optimization of operational networks. One of the most important goals of traffic engineering is to achieve efficient use of network resources, particularly the available bandwidth. In today's IP networks, traffic flows are mostly not optimally mapped to the available resources. The main reason for this is that in this traditional routing architecture with the deployed routing protocols there is no straightforward methodology for performing traffic engineering; the paths of the traffic can only be influenced implicitly by changing the setting of link weights in the network, which makes traffic allocation in an intuitive manner almost impossible. New approaches and technologies for traffic engineering in the Internet are therefore required. In addition to traffic engineering, the optimization of IP networks with respect to resilience requirements is becoming an increasingly important research topic. Usually, the goal of such optimization efforts is to ensure non-degraded levels of service in the presence of single link or node failures with a minimum of extra capacity required for resilience.

In the framework of the COST 279 Action, a variety of approaches concerning research methodology and employed technologies has been investigated. The individual proposals are introduced in the next subsections.

1.5.1 Traffic Engineering Approaches Based on IP Routing

Recently, many competing algorithms have been proposed which enhance the current routing algorithms and protocols with traffic engineering capabilities. These algorithms often employ multipath routing because it opens the possibility of achieving an even distribution of load in the network. However, the standard Equal Cost Multi-Path (ECMP) routing usually does not provide a sufficiently large number of multiple paths, as all paths are required to have the same, minimal cost. Additionally, ECMP always distributes traffic evenly among the available paths without taking into consideration the current traffic conditions in the network.

Multipath Routing with Dynamic Variance (MRDV) is proposed in [49] as a routing protocol which improves the performance of IP routing protocols under overload conditions. The main objective was to design a scalable and

simple dynamic routing algorithm compatible with the traditional IP protocols. In the MRDV scheme, whenever a node tries to send too much traffic through a path, this traffic is distributed among several additional paths. A new routing metric called “multipath with variance” is proposed, which enables the traffic for each destination to be carried by additional paths beside those with minimal cost. The “variance” parameter basically describes the extent to which the cost of additional paths may exceed the cost of optimal paths. The central concept of MRDV is that the number of such additional paths towards a destination dynamically depends on the extent of link overload. This multipath scheme results in a reduction of congestion and leads to a better use of network resources. The MDRV algorithm has been evaluated using Network Simulator (ns-2) for several different scenarios. The conducted performance evaluation shows the algorithm’s stability as well as considerable performance improvements over traditional routing schemes.

In [50], Adaptive Multi-Path Routing (AMP) is proposed as a new routing algorithm for dynamic traffic engineering. This distributed algorithm is envisioned as an extension to current intra-domain routing protocols, which operates autonomously and continuously in the nodes of the network, such that it does not introduce any management overhead. With AMP, every node performs load measurements on its output links. In the case of congestion, the node which is sending traffic out directly on the congested link will immediately react to overload by trying to put more load on available alternative paths. At the same time, this node will send so called backpressure messages to each of its neighbor nodes, notifying them to which extent traffic they are sending is present on the congested link. In other words, the neighbor nodes are informed about their contributions to congestion. These neighbors will then also try to offload their links which lead towards the congestion hotspot and they will pass similar congestion indications (i.e., backpressure messages) to their neighbor nodes, etc. This quasi-recursive signaling mechanism of backpressure messages is the central innovation of AMP. It enables global propagation of load information, and at the same time it keeps the exchange of load information (i.e., the signaling) local and thus very efficient with respect to the consumption of network resources. Having a large number of multiple paths in the network is a prerequisite for efficient load balancing, and therefore it is important to also consider additional paths beside those with minimal cost. In order to use such non-minimal paths, and still strictly avoid routing loops, AMP uses a

simple criterion which states that any neighbor node which is closer to the destination in terms of cost than the current node is a viable next hop for multipath routing. Note as an immediate conclusion that routing loops cannot form as the cost to reach the destination always strictly decreases along every hop of the path. AMP uses the 16-bit Cyclic Redundancy Check (CRC-16) for splitting traffic into microflow aggregates based on the {source, destination} address pairs, and divides the traffic among multiple paths by mapping portions of the CRC-16 hash space (i.e., a range of microflow aggregates) to the viable next hops for each destination. Load balancing is performed such that the relative sizes of the hash space portions for each destination are dynamically adjusted in every node. The algorithm has been implemented in Network Simulator (ns-2), and simulations of the 27 node and 47 links AT&T-US network have been carried out using a state of the art Web traffic model. The performance of AMP is compared to shortest path routing and equal cost multipath routing for a broad spectrum of load. The results demonstrate significant performance improvements and the stability of AMP throughout the investigated simulations.

It is well known that mixed integer programming must be used in order to determine the optimal set of link weights for an IP network. This makes the optimization problem NP-complete, such that heuristic method must be used, which may result in inappropriately long computation times for large networks. Therefore, a clear motivation is given for research in reducing the scalability limitations of IP routing optimization. A novel method for the optimization of large IP networks called “divide and conquer” is presented in [51]. The basic idea is to split the original network into smaller subnetworks and to optimize these smaller subnetworks individually. In the “divide phase” of the algorithm, the network which is to be optimized is decomposed into three subnetworks. Two of these subnetworks are disjoint and relatively large, and the third subnetwork is designed as a small central “link” between them. Once the decomposition of the original network is accomplished, the “conquer phase” may begin. In this phase mixed integer programming is used for the formulation of the optimization problem for the two large subnetworks created in the “divide phase”. In the final step of the algorithm, after the optimization of these subnetworks has been completed, it is left to merge the two individual results into one solution for the entire original network. Both the performance of the decomposition algorithm itself and the overall routing optimization using decomposition have been evaluated. The performance of the overall “divide and

conquer” scheme has been compared to both the direct optimization results for the original network, and to the non-optimized original network (the topology of the 25 node AT&T-US and a smaller 14-node network were used). Results obtained from the different optimization techniques lead to the conclusion that although “divide and conquer” never reaches the performance of direct optimization, it represents a valuable technique for the optimization of large networks due to the significant reductions of computation times it achieves.

1.5.2 Traffic Engineering and Resilience Approaches Based on Multi-Protocol Label Switching

As already elaborated in detail in the previous subsection, new distributed algorithms for traffic engineering are introduced in [50] and [49]. Their common characteristic is that no additional management overhead (i.e., no additional intervention by network operators) is required for performing traffic engineering. However, both algorithms also require the enhancement of the network (i.e., its nodes) with the specific functionalities required. In [52, 53, 54, 55], a different approach is preferred where traffic engineering is performed by network management based on Multi-Protocol Label Switching (MPLS). MPLS is a technology for establishing LSPs between pairs of nodes in a network domain. Unlike what happens in the traditional IP architecture, where the paths of traffic are derived from the assigned link weights, the paths can be chosen arbitrarily with MPLS. This opens a large potential for traffic engineering and for novel resilience schemes, provided that a comprehensive concept is used for path establishment and network management. As no additional networking algorithms or protocols are generally needed if MPLS is used, no updates to the network technology are required. The essential mechanism in most of these approaches is to use offline optimization of LSPs. Traffic engineering methods for MPLS networks are in many ways similar to methods developed for ATM (like, e.g., in [55]), as both technologies essentially rely on establishing label switched paths in the network.

Several different offline traffic engineering approaches for MPLS networks are presented in [52]. In the first approach, only MPLS is used for the routing of traffic, and traffic engineering is performed by defining a linear program for network optimization. The input for the program is information about the network topology and the traffic demand matrix. The weighted sum of the av-

erage and the maximum link utilization is minimized in the objective function. For the purpose of performance evaluation, the weights are set such that either the maximum, or the average, or both the maximum and the average link utilization are minimized. Different optimization strategies are evaluated for two different networks of 20 and 25 nodes. The values for the average and the maximum link utilization obtained in these optimizations are compared to values which result from using pure IP routing. As expected, the best results in MPLS optimization are achieved when both the average and the maximum link utilization are included in the objective function. In the second approach three different routing scenarios are compared (unoptimized pure IP network, IP network with optimized link weights, and optimized MPLS network) for networks of 10, 20, and 25 nodes. The results show that the optimized MPLS and the optimized pure IP network perform similarly well (with a slight advantage for MPLS), and that both optimized schemes by far outperform the unoptimized pure IP network. As MPLS is able to operate simultaneously with the IP protocol in the same network without interferences, traffic engineering possibilities for a combination of IP and MPLS traffic forwarding are also investigated. A new algorithm for traffic engineering of combined pure IP and MPLS networks called Decompose-Design-Reassemble (DDR) is introduced. The input for the algorithm is a pure IP network in which a number of flows are identified as suitable for traffic engineering manipulations. Subsequently, LSPs which are supposed to accommodate these flows are computed. These LSPs will be established only in the case that the computed path of the LSP is not identical with the corresponding IP path for a particular flow. This results in a significant reduction of the number of flows which will actually be forwarded within LSPs. Results show that usually only up to one third of the paths computed by such an optimization algorithm must really be routed within LSPs because they do not match the standard IP paths. This opens the potential for significant reductions of the amount of MPLS-related state information in the network.

In contrast to traffic engineering, which is usually performed on the time-scale of minutes or hours, network capacity planning is a long-term strategic activity for network operators. Besides forecasting future traffic demands, a major factor which must be considered in the context of planning is the amount of backup capacity needed for ensuring resilience. As the capacity installed for resilience will not be used under normal network conditions (i.e., in the ab-

sence of network faults), network operators have a strong economic incentive to minimize the amount of backup capacity installed. In [30], new protection switching mechanisms for autonomous systems are proposed which may be implemented using MPLS. The fundamental idea is to take advantage of the load balancing potential of multipath traffic forwarding in order to minimize the amount of required backup capacity. Inside the autonomous system, only disjoint multiple border-to-border paths are employed, such that network configuration and failure diagnostics are kept simple. Two basic path protection mechanisms are proposed, which differ with respect to the use of multiple paths. In the first approach, traffic is routed only on single shortest paths in the absence of failures, and multipaths are used for backup purposes. The second mechanism employs multiple paths, and in the case of failure it redistributes traffic from the inactive paths to the active paths. The results show that the performance of the introduced protection schemes depends on the network topology, and in particular on the number of disjoint paths in the network, but not on the network size. The main metric used in the performance evaluation of the presented mechanisms is the backup capacity required. It is demonstrated that the proposed schemes require significantly less backup capacity than, e.g., standard Open Shortest Path First (OSPF) re-routing. The main result of the performance evaluation is that typically only around 20% additional resources are needed to provide full resilience against all single link or node failures in well designed networks.

1.6 Scheduling

This section reports a model of a specific scheduler, theoretical results about instability phenomena in packet networks, a new priority scheme for the QoS provisioning in the Internet, and the use of traffic prediction for dynamic resource allocation schemes.

A Lambda scheduler within a Generalized MPLS (GMPLS) node is studied in [56] both by means of an OPNET simulation model and by an analytical model. The simulation model consists of two basic elements: the request generator and the λ -scheduler. The request generator describes the arrival process of the requests that are generated for a GMPLS node by its adjacent nodes. The user can configure the simulator so that the inter-arrival time of requests matches a desired statistical distribution. The λ -scheduler serves the requests

according to the following wavelength mapping policy. First, the scheduler checks if an output wavelength is available (a constraint, i.e., a list of possible output wavelengths, can be associated to each request). Then, in order to save wavelength converters, preference is given to the output wavelength which directly maps the input one. If direct mapping is not possible, the availability of a wavelength converter is checked. If no converters are free, the request is rejected. For the development of the analytical model, which is based on a Markov chain, the arrival process of requests is assumed to be Poisson and the outlets are chosen at random. Requests are lost if no direct mapping and no converter are available. The accuracy of the analytical model is validated by comparison against simulation results. Moreover, the performance of the scheduler is evaluated for different values of node size, number of converters, and offered load. Performance is evaluated in terms of the probability to reject a call and of the node utilization. Clearly, the performance degrades with the traffic load and improves with the size of the node and the number of wavelength converters.

Instability phenomena in underloaded packet networks are investigated in [57]. Instability may occur in underloaded networks under particular conditions, such as routes that make the customers visit the same queues several times, variations of the customer service times at different queues, or complex scheduling algorithms. In [57], QoS schedulers in packet networks are considered. In order to represent the approaches currently considered for QoS provisioning in the Internet (e.g., DiffServ), the analyzed scenarios have acyclic routes and service times varying according to the channel capacity. Two kinds of schedulers are considered: Generalized Processor Sharing (GPS) and strict priority. Networks of queues with GPS schedulers were already known to be universally stable provided that nominal rates of the flows are greater or equal than the effective average rates. As a new result, in [57] it is shown that some instability can occur when some of the actual packet rates exceeds the nominal rates, which is quite likely to occur in real systems. Indeed, a mismatch between nominal rates and effective rates is likely since the estimation of the effective rates is typically based on local traffic measurements and tends, thus, to be inaccurate. For what concerns strict priority schedulers, they are proved to be stable provided that the priority ordering of packet flows is the same in all the queues of the network. Moreover, weak forms of instability can be observed in networks which combine First In First Out (FIFO) and priority

scheduling for values of the load as low as 0.6.

A new method called Priority Forcing Scheme (PFS) is analyzed in [58]. PFS allows Internet applications to get service better than best effort by means of the following mechanism. Besides data packets, the PFS application generates some additional redundant packets, named reservation packets, whose role is to implicitly reserve resources for a given data traffic flow. Indeed, at its arrival at a router, a data packet checks if any reservation packet belonging to the same flow is already queued in the router. If so, the data packet can replace the reservation packet, so that the data flow receives a prioritized service. Numerical results show that the performance improvement achieved by PFS increases with the average number of queued packets, i.e., as the traffic intensity increases. By properly setting the parameters of the scheme, packet delay characteristics can be shaped. As cases of practical interest, Voice Over IP and FTP applications are considered. The proposed scheme improves the end-to-end quality of both services, and is particularly efficient under overload traffic conditions.

The authors of [59] consider a link shared between high priority traffic and some low-priority streams. In order to optimize the bandwidth allocation, the algorithm which decides how to share the available bandwidth between the low-priority flows can use estimations of the future load. The document provides a framework for evaluating the possible benefits that dynamic resource management algorithms achieve by employing traffic prediction techniques. The performance of a processor sharing system is evaluated by simulation. In the considered experiments, the traffic prediction technique is based on adaptive filters, which are capable of properly responding to traffic variations. Moreover, the objective of the resource allocation scheme is to keep the input queues of the same size.

1.7 Multicast

The popularity of large-scale distributed applications, such as video conference, multimedia dissemination, electronic stock exchange, and distributed cooperative work has grown with the availability of high-speed networks and the expansion of the Internet. The key property of these types of applications is the need to distribute data among multiple participants together with application specific QoS requirements. This fact makes scalable multicast protocols an

essential underlying communication structure. Increase in the share of multicasting in high-speed data networks is expected to alter the aggregate network traffic structure and call for improved control mechanisms. Investigation regarding the influence of QoS parameters such as reliability and scalability on network traffic will be important.

A significant protocol element for ensuring reliability at the transport level is message loss recovery. Two representative approaches for many existing loss recovery solutions in scalable multicasting are nonhierarchical feedback control and hierarchical feedback control. Overall, the key issue is to reduce the number of feedback messages that are returned to the sender. In the former approach, a model that has been adopted by several wide-area applications is referred to as feedback suppression. A well-known example is the Scalable Reliable Multicast (ScRM) protocol. In ScRM, when a receiver detects that it missed a message, it multicasts its feedback to the rest of the group. Multicasting feedback allows another group to suppress its own feedback. A receiver lacking a message schedules a feedback with some random delay. An improvement to enhance scalability is referred to as local recovery, which is related to restraining the recovery of a message loss to the region where the loss has occurred. In the latter approach, hierarchical approaches are adopted for achieving scalability for very large groups of receivers. Scalable reliable multicast is an area of active research. In addition to the mentioned approaches, another alternative for ensuring reliability is Forward Error Correction (FEC). The idea behind this approach is predicting losses and transmitting redundant data. On the other hand, recent epidemic or randomized approaches to loss recovery have promising outcomes in terms of robustness and overhead. In this direction, Bimodal Multicast provides an epidemic loss recovery mechanism. Periodically, every site chooses another site at random and exchanges information to see any differences and achieve consistency. The information exchanged includes some form of message history of the group members. Epidemic protocols are simple, scale well, are robust against common failures, and provide eventual consistency as well. They combine benefits of efficiency in hierarchical data dissemination with robustness in flooding protocols.

Both [60] and [61] focus on the traffic properties of Bimodal Multicast induced by its epidemic recovery mechanism in comparison to ScRM. The empirical results demonstrate that Bimodal Multicast generates more desirable traffic than ScRM. We elaborate on the protocol mechanisms as the underly-

ing factor in these results. Further description can be found in Section 2.3.1. From a traffic perspective, the aim of these studies is to discover and develop multicast protocols that not only feed well-behaved traffic discretely into the existing networks, but also can cope with the existing self-similar traffic and its adverse consequences.

A large deviations approximation for evaluating the gain of using multicast over unicast in a communications link is presented in [62]. This is a useful approach, especially for large links, as Monte Carlo simulation has problems in this case. Multicast gain is defined as the ratio of number of users that can be served with a target blocking probability with multicast divided by the number of such users with unicast. A simpler estimate of multicast gain is the ratio of the number of users that generate mean link occupancy in each communication method. The idea in the combination of unicast and broadcast is to use broadcast to transmit the most popular content. For this to work, the operator needs to select the appropriate content and allocate the optimal number of channels that minimizes the blocking probability of unicast. Gain over unicast is computed similarly with a combination of unicast and broadcast instead of multicast. Numerical results show that for larger values of the parameter of the Zipf distribution, valid for a user's selection of a particular channel, the combination of unicast and broadcast outperforms the use of multicast. The results contain approximations and rely on the assumption that the operator is able to select the most popular channels. However, broadcasting may be considered as an alternative also due to the concerns about practical problems associated with implementing the multicast functionality in the network.

[63] focuses on the epidemic anti-entropy model of Bimodal Multicast utilized for reliability of the protocol, and demonstrates its scalability and robustness. We give our comparative simulation results discussing the performance of the model on a range of typical scenarios. We demonstrate that anti-entropy loss recovery produces balanced overhead distribution among group receivers and is scalable with an increase in group size, multicast message rate, and system-wide noise rate. The anti-entropy distribution model provides eventual data consistency, scalability, and independence from communication topologies. It also transparently handles failures and can work with minimal connectivity.

Overall, the results of the above studies can be considered toward the general problem of integration of multicast communication into the Internet. In

particular, mechanisms including efficient loss recovery and scalability to support multicast transport would be essential. Various applications on both wired and wireless settings could benefit from multicast communication support with efficient provision of reliability guarantees.

1.8 Graph Models for Interconnection Networks

The problem of finding an adequate characterization of the topology of the huge, global Internet has been an object of intensive research during the last few years. The necessary prerequisite for such studies is provided by the empirical results on Internet topology, obtained from systematic large-scale sensing with `traceroute`. The Cooperative Association for Internet Data Analysis (CAIDA; www.caida.org) has been most important in this activity.

One of the most interesting models studied in this context is the so called Power-Law Random Graph (PLRG). The degrees of the nodes are drawn independently from the power-law distribution, and the connections are then chosen randomly. This model is extremely simple, and it has been a surprise that its resemblance to the real Internet is much better than one might think. Three COST 279 contributions have been devoted to this model. (It should be noted that these papers focus on the model itself, not on its relation to reality.)

The main result in papers [64] and [65] is the following: Denote the number of nodes by N and assume that the tail of the degree distribution decreases as $P(D > d) \sim d^{-\tau+1}$, where $\tau \in (2, 3)$ (infinite variance!). Then the hop distance between two randomly selected nodes of the giant component (the random graph need not be fully connected) is asymptotically almost surely less than $(2 + o(1))k^*$, where

$$k^* = \left\lceil \frac{\log \log N}{-\log(\tau - 2)} \right\rceil.$$

Thus, the practical diameter of the graph grows hardly at all when N runs through many orders of magnitude. This phenomenon, and the proof of the theorem, is based on a kind of “soft hierarchy”: the nodes with high degrees (“high” meaning larger than a certain function of N) form a “core network” that enables short connections.

The paper [66] makes further studies on the model, mainly based on simulations only. A routing system is added to the core part of the model by assuming a “natural” routing system, where the routing is hereditary, invertible, and uses least-hop paths. This lets the traffic concentrate very unevenly on the links. With a uniform traffic matrix, a roughly exponential distribution of link loads was observed in this setup.

Since the nodes with large (high degree) nodes play a central role in the high connectivity of the PLRG, it was interesting to study the vulnerability of the PLRG against losses of large nodes. When only the huge nodes with degree higher than \sqrt{N} were removed, the distances became a little longer, but the effect was not dramatic. When the whole core was lost, a big part of the network remained still connected, but now the distances became substantially longer. These results tell of very high resilience of the topology but are probably overoptimistic in comparison with the real Internet.

The last topic studied in [66] were the implications of the model to the efficiency of multicasting. It was found that, in a moderate size multicast tree between randomly selected nodes, the gain is obtained mainly in the portion from the source to the center part of the core, after which the routes to the different receivers don’t overlap much. This differs rather clearly from the more standard scenario of branching resembling a binary tree.

Chapter 2

Traffic Measurement and Characterization

Markus Fiedler

Blekinge Institute of Technology, Karlskrona, Sweden

Contributors:

Mine Caglar (Koc University, Turkey), Patrik Carlsson (Blekinge Institute of Technology, Sweden), Marcus Fiedler (Blekinge Institute of Technology, Sweden), Guoqiang Hu (University of Stuttgart, Germany), Jorma Kilpi (VTT Information Technology, Finland), Udo Krieger (Otto Friedrich University Bamberg, Germany), Petteri Mannersalo (VTT Information Technology, Finland), Sándor Molnár (Budapest University of Technology and Economics, Hungary), Luca Muscariello (Politecnico di Torino, Italy), Philippe Olivier (France Telecom R&D, France), Öznur Özkasap (Koc University, Turkey), Adrian Popescu (Blekinge Institute of Technology, Sweden), Marie-Ange Remiche (Université Libre de Bruxelles, Belgium), Kavé Salamatian (Université Paris VI, France), Paulo Salvador (Telecommunications Institute, Portugal), Kurt Tutschku (University of Würzburg, Germany), Hans van den Berg (TNO Telecom, Netherlands), Sabine Wittevrongel (Ghent University, Belgium).

2.1 Introduction

We observe currently a fast evolution of web, multimedia, and especially peer-to-peer applications. This evolution yields a rapidly changing load environment for the transport of data streams subject to severe QoS constraints. Considering the needs for traffic planning and QoS assessment in wired and wireless access, as well as backbone networks, issues like load modeling and traffic characterization by thorough statistical techniques have to be addressed as the first basic steps of teletraffic engineering. For this purpose, new and reliable traffic measurement, modeling, and estimation techniques are required to ac-

curately determine the traffic load and the resource usage in wired and wireless networks. The goal of this chapter consists in providing a comprehensive overview on related activities within the COST 279 action. Its structure reflects natural steps of obtaining results:

- *traffic measurement* in Section 2.2, which concerns rather technical issues related to the collection of traffic and its statistics;
- *traffic characterization* in Section 2.3, to provide a rather qualitative description of traffic behavior;
- *traffic modeling* in Section 2.4), to provide a rather quantitative description with focus on models and their parameters;
- *traffic estimation* in Section 2.5), which is about inference and prediction of traffic volumes based on observations.

2.2 Traffic Measurement

The process of measuring traffic forms the basis not only for characterization, modeling, and estimation, but also for network dimensioning, traffic engineering, capacity provisioning, and management. By providing an important link to reality in order to establish and validate methods and actions, traffic measurements play a central role for almost any kind of network-related activity. However, measurements in the Internet context have traditionally not been carried out with that much effort and precision as in the case of telecommunications, where they take a very natural place in QoS provisioning.

The following subsections provide an overview of tools used in various studies and discuss achievements in measurement methodology developed in the scope of COST 279.

2.2.1 Passive Measurements

[67] defines: “Passive measurements are by nature non-intrusive. In this class of measurements, traffic parameters are monitored at a particular point of the network such as a router or a Point of Presence (POP).” By monitoring (i.e., basically counting) or collecting traffic in different places at the same time, the analyst may get a holistic view on what is going on in a network and collect statistical data to build models upon.

For network operation and management purpose there exists a widely deployed infrastructure for monitoring performance-related data: almost each non-low-cost IP-enabled device offers an agent that can be polled via the Simple Network Management Protocol (SNMP). Sometimes, the Remote Monitoring (RMON) extension delivers additional performance data. For performance monitoring purposes, such a polling is done on rather large time scales (5–60 minutes). As advanced traffic modeling needs information about the traffic on timescales of seconds and beyond, [68] investigates to which extent SNMP can be used for such a purpose. Depending on the device under test, counter update intervals of several seconds or response times of up to half a second were observed. This illustrates the need of evaluating the response behavior of SNMP agents before using them for quantitative measurements. In case of backbone links, successful efforts to infer traffic behaviour on short time scales from medium-term measurements (e.g., 15 minute intervals) are presented in [36].

If information more detailed and accurate than the one provided by SNMP and RMON is required, packets (or at least their headers) need to be collected with the corresponding time stamps. A very common tool for passive monitoring of IP packets is `tcpdump`, which is used in [69, 70, 71, 72, 73, 74, 75, 76]. Usually the tool runs on a sender, receiver, or dedicated monitoring machine, for instance connected to a monitoring port of a switch [71] (see Figure 2.1), or to a wiretap. `Tcpdump` traces have to be post-analyzed to yield the desired information. For example, the tool `Tstat` [75] allows for the reconstruction of TCP sessions, while [76] describes how to identify signalling and download packets for the currently very popular peer-to-peer service eDonkey [77].

There are however critical issues related to the synchronization and precision of the time stamps [67, 78]. Moreover, as the measurement equipment is usually expensive, there is an economical gain when sharing measurement equipment between measurement processes. A concept of distributed Passive Measurements Infrastructure (PMI) is presented in [78]. This paper discusses on how to coordinate and manage such joint measurement equipment. It does so using the concepts of Measurement Points (MPs), Measurement Areas (MAs), and Super Measurement Areas (SMAs). To allow for large scale deployment the PMI can use a wide range of collection equipment, from dedicated hardware to Packet CAPture (PCAP) library solutions. Furthermore, the PMI focuses on capturing and collection, while the analysis of the data is up




Figure 2.1: Example for the use of `tcpdump` on a monitoring port of a switch [71].

to the user.

Yet another possibility of implementing passive measurements consists in adding the desired functionality to (public-domain) software. For instance, in [79] a Gnutella servent process `Gnut` was modified in order to collect the desired statistics.

Finally, [9] describes the on-line optimization of token bucket parameters based on on-line measurements of the target stream.

2.2.2 Active Measurements

According to [67], “Active measurements are more intrusive, as they inject traffic into the network. The rationale behind active measurements is that estimating the end to end QoS, as sensed by a real application, can only be done by putting oneself in place of the real application.” Thus, active measurements reflect the experience of a network user rather than a holistic view of the overall performance of a network.

The simplest and most widely known active measurement tool is `ping`, which allows for the estimations of RTT and loss quotas.[80] presents a further development for wireless networks called *Wave-Ping* (WVPing). When invoked, this tool reads status information from the wireless network interface card driver. Thus, the important connection between RTT and signal quality, noise level, Signal-to-Noise Ratio (SNR), and number of erroneous packets on Medium Access Control (MAC) level can be displayed. Furthermore, the number of the radio channel reveals handovers. WVPing results are shown in Section 2.3.6.

Another example of an active measurement is given by the *packet-pair technique* described in [67]. Two packets of length L are sent into the network in a back-to-back fashion. If they arrive dispersed at the receiver, i.e., spaced by time τ , they have experienced a bottleneck of capacity $\mu = L/\tau$.

2.3 Traffic Characterization

The following subsections contain a broad realm of qualitative studies of traffic behaviors in different scenarios. Multicast, Optical Burst Switching, and TCP traffic are investigated with regard to self-similar properties. Also, traffic aggregation for modem pool traffic yielding Gaussian properties is discussed. Finally, General Packet Radio Service (GPRS), Wireless LAN (WLAN), and Peer-to-Peer (P2P) traffic are characterized.

2.3.1 Multicast traffic

Most efforts on multicast communication have focused on developing new applications and protocols that are compared with respect to performance measures such as scalability, reliability and congestion control. However, the nature of the traffic stream generated by each type of protocol, particularly with respect to self-similarity, has not been studied extensively. The transport layer mechanisms are important components in translating heavy-tailed file size distributions at the application layer into link level traffic self-similarity. In particular, [60] and [61] focus on the traffic that scalable multicast protocols generate. The results provide empirical evidence for the fact that the level of self-similarity depends on the transport protocol used. Bimodal Multicast is a novel option in the spectrum of multicast protocols. It is based on an epidemic loss

recovery mechanism which has a Markovian structure. The protocol has been shown to exhibit stable throughput under failure scenarios that are common on real large-scale networks. In contrast, this kind of behavior can cause other reliable multicast protocols to exhibit unstable throughput. Bimodal Multicast is compared to ScRM as the latter is inspired by the principles of TCP/IP protocol stack, which is prevalent in the Internet and exhibits self-similarity.

In [60], empirical results demonstrate that the epidemic approach of Bimodal Multicast generates more desirable traffic than ScRM in the case of a Constant Bit Rate (CBR) source. ScRM traffic shows long-range dependence and self-similarity whereas Bimodal Multicast traffic is short-range dependent. The approach is to analyze simulation traces obtained from ns-2 and provide analytical support for the results. The delays at the transport level and traffic at the link level are studied. Bimodal Multicast yields lower overhead traffic and transport delays than ScRM. The protocol mechanisms are elaborated on as the underlying factor in the empirical results. The intrinsic relation of these mechanisms to traffic characteristics is studied through delay analysis in the case of Bimodal Multicast and discussed for ScRM. The marginal delay distribution has been analyzed for Bimodal Multicast and is shown to decay exponentially. This substantiates that Long Range Dependence (LRD) is not expected intrinsically due to the epidemic mechanism of the protocol. More information on epidemic loss recovery can be found in Section 1.7.

[61] considers a self-similar source, namely an on/off sender that transmits with Pareto on and off times. It is well known that when sufficiently many of traffic streams from such sources are aggregated, long-range dependence arises at the link level. Results pertaining to self-similar source are compared with the CBR case. In the case of an on/off sender, Bimodal Multicast still generates short-range dependent delay, which is scalable in number of users like its other superior performance properties of stable throughput, lower protocol overhead, and reliability. However, long-range dependence emerges at the link level when the source induces self-similarity. Although there is a single on/off source, at the link level there is an aggregation arising from the recovery process of all receivers in the network. For ScRM, the traffic becomes worse in terms of both delays and at the link level compared to the CBR case. The timer-based loss recovery mechanism of ScRM triggers self-similarity with long-range dependence, whereas the Markovian structure of Bimodal Multicast inherently generates short-range dependent traffic. As future work, Bi-

modal Multicast deserves more research in terms of traffic characteristics for various parameters, such as gossip rate and buffer sizes. One goal for such work is to provide a stochastic model involving the parameters and the mechanisms of Bimodal multicast. Comparative studies with other scalable multicast approaches will help to identify efficient protocol mechanisms.

2.3.2 Optical Burst Switching traffic

In Optical Burst Switching (OBS) networks, the essential function of *burst assembly* puts together, at the edge nodes, a number of IP packets into an optical burst of larger size so as to achieve efficient burst switching. The ability of OBS to reduce the self-similarity of IP traffic, heavily discussed recently, is important for the evaluation of OBS network performance. In [81] this question is studied both via simulation and via theoretical analysis. Three different assembly schemes are inspected: threshold-based scheme, time-based schemes, and time-based with padding. The traffic is studied in terms of byte and packet (or burst) stream, respectively. It is found that only in the case of time-based assembly and regarding the burst departure process, self-similarity is reduced when increasing the time-out interval of the assembly algorithm. In all other cases, especially when traffic is measured as byte-stream, assembly has no impact on the self-similarity properties.

2.3.3 TCP traffic

[70, 82] present a measurement and analysis study on how TCP congestion control can propagate self-similarity between distant areas of the Internet. This property of TCP is due to its congestion control algorithm, which adapts to self-similar fluctuations on several timescales. The mechanisms and limitations of this propagation are investigated, and it is demonstrated that if a TCP connection shares a bottleneck link with a self-similar background traffic flow, it propagates the correlation structure of the background traffic flow above a characteristic timescale. The cut-off timescale depends on the end-to-end path properties, e.g., round-trip time and average window size. It is also demonstrated that even short TCP connections can propagate long-range correlations effectively. The analysis reveals that if congestion periods in a connection's hops are long-range dependent, then the end-user perceived end-to-end traffic

is also long-range dependent and is characterized by the largest Hurst exponent. Furthermore, it is shown that self-similarity of one TCP stream can be passed on to other TCP streams multiplexed together with it. These mechanisms complement the widespread scaling phenomena reported in a number of recent papers. The arguments are supported with a combination of analytic techniques, simulations and statistical analyses of real Internet traffic measurements.

2.3.4 Modem pool traffic

Gaussian models are discussed in paper [83] and in its preliminary version [69]. Whereas [69] is more oriented to traffic characterization, the final version [83] turns more to the methodological direction. Any Gaussian traffic model is motivated by the Central Limit Theorems (CLTs). Hence, to justify Gaussian models there must be enough traffic aggregation. The aggregation can be theoretically divided into the *horizontal* and *vertical* directions. Horizontal aggregation is the time scale or resolution, aggregation is essentially the number of contributing sources at a time slot.

The traffic trace analyzed in [69] and in [83] was measured from a modem pool of a commercial Internet Service Provider (ISP) and, since the level of vertical aggregation was rather high, a natural question to ask was whether it was high enough to justify Gaussian traffic models for some (or any) time resolutions. [69] contains already the basic ideas of determining the minimum time scale and minimum number of contributing sources that would be required by the CLT. In [83] some simple and pragmatic ideas were used to determine the first time scale where the Gaussian approximation is plausible. In the downstream directed traffic this was between 0.512 s and 1.024 s for the data trace studied. However, the effect of the CLT was already visible from a 0.128 s resolution, but the amount of vertical aggregation was not sufficient for resolutions from 0.128 s to 0.512 s. Some heuristic approximations were made to determine how much more vertical aggregation would be required in these smaller time scales in an ideal case. In the upstream directed traffic none of the timescales studied gave satisfactory results, and some reasons for this are also explained in [83].

2.3.5 GPRS traffic

GPRS traffic is considered in [84] and in its preliminary version [72]. These papers present some results obtained from `tcpdump` traces recorded between the GPRS backbone network and the Internet, but before the Network/Private Address Translation (NAT) is done. The measurement was done in May, 2002, about half a year after the operator had launched its commercial GPRS service. It is probably one of the first published GPRS user traffic measurements in the world. A GPRS session was defined as all the packets (and their time stamps) with the same temporary IP address that the user is given when he/she is attached to the GPRS network. The main observations were that during the measurement time, most GPRS sessions started during the working days and in working hours, that they were in general roughly similar to low access speed dial-up sessions, and that GPRS session durations and session volumes seem to have heavy-tailed distributions. Also, the majority of the data transfer occurs typically in the beginning of the session and the GPRS users typically detach from GPRS network when they have finished active usage.

[74] studies distributional properties of GPRS/GSM session volumes and durations. It combines a known statistical method of ascertaining about the underlying distribution with the data-analysis of GPRS data introduced in [72] and in [84]. The statistical method is based on the concept of maximum correlation, and was used because there was no *a priori* reason to expect or prefer any particular distribution. The main result is that these data sets exhibit a heavy-tailed nature. The Weibull distribution with the shape parameter between $0 < \alpha < 2$ is one of the possible models. [74] is a first draft version of the document; please contact the author of the paper for the latest version and the current status of the document.

2.3.6 WLAN traffic

The integration of wireless networks into high-speed NGNs based on the IP technology and its corresponding resource reservation and QoS-management mechanisms is one of the most challenging current tasks of network design and engineering. The fast convergence of these wireless and wired networks coincides with the planned deployment of new integrated multi-media and interactive Web services that are independent of the realized mobility patterns of the terminals and users. Traffic characterization in such wireless environments




Figure 2.2: RTT delay samples from WVPing with 4088 byte payload for roaming in a basic service area of the 2 Mbps environment.

constitutes a new exciting task of teletraffic theory.

In [80] a wireless next generation IP-network and its corresponding QoS-management mechanisms are considered. A service architecture is sketched where adaptive applications are running on top of the transport and network layers with their resource and mobility-management functionalities. A series of comprehensive traffic measurements is used to study the interworking of flow control, resource reservation and mobility management. Considering a terminal roaming in a basic service area of a 2 Mbps IEEE802.11 WLAN, typical results of the RTT as compared to the SNR and related handovers, as well as the evolution of the TCP congestion window, are shown in Figure 2.2 and 2.3, respectively.

2.3.7 P2P traffic

P2P services have evolved to one the most popular applications in today's Internet. P2P networks have become very popular recently, as witnessed by the relentless spread of the Gnutella [85], Kazaa [86], and eDonkey [77] file sharing applications. Remarkably, only very simple protocols and almost no support by the transport network have been needed to make these distributed




Figure 2.3: Dynamics of the TCP congestion window subject to roaming in one basic service area for stream traffic patterns in the 2 Mbps environment.

services operable on a large scale in very little time. P2P services trade their advantages of autonomous, adaptive, and resilient service operation with the disadvantage of causing high data and signalling traffic volumes. The traffic patterns of P2P applications fluctuate strongly in time and space. As a result, it is anticipated that traditional network design techniques and traffic engineering procedures may no longer be applicable and new methods may be needed that maintain the autonomous and self-organizing characteristics of P2P.

In [79], the authors present a measurement study on the signalling traffic in Gnutella overlay networks. Both signalling load and the scale of variability in the existence of P2P overlay connections are investigated. The study reveals that an uncontrolled Gnutella client consumes high amounts of bandwidth for its signalling traffic, reaching up in the order of tens of Mbps. Furthermore, the signalling traffic in Gnutella overlays varies strongly over short timescales, due to the Gnutella use of flooding protocols. The investigation of the overlay connection holding time in Gnutella showed that the distribution typically has bi-modal characteristic. The modes correspond to a “short” state, where typically host information is transmitted, and to a “stable” mode, where mainly content queries are exchanged. The modes identify the time scales on which a dynamic and adaptive management of P2P overlays and P2P services is of advantage or needed.

In [76] the authors provide a traffic profile for the eDonkey P2P filesharing




Figure 2.4: Correlation of eDonkey TCP holding time and flow size

service [77]. The eDonkey system is typically used for exchanging very large files like audio/video CDs or even DVD images, and possesses a hybrid P2P architecture with distinct servers and clients. The eDonkey system makes use of the multi source download feature, which permits the simultaneous transmission of file chunks to a downloading peer. The traffic profile shows that signalling and download have significantly different characteristics. Figure 2.4 depicts a scatter plot describing graphically the correlation of the TCP holding time and the size of eDonkey flows. Each dot in the scatter plot represents an observed eDonkey flow. The brighter dots are identified download flows, the dark dots represent non-download connections. The scatter plot shows that almost all identified download flows are within the same region. In turn, the non-download flows are in an disjunct region of the plot. This graph reveals that download and non-download flows have significantly different characteristics. A future traffic model has to distinguish between both types of traffic. In addition, the traffic profile, provided in [76], gives evidence that the expected “mice and elephant” phenomenon in eDonkey traffic is not as severe as expected.

2.4 Traffic Modeling

This section starts with contributions on the estimation of probability density functions (pdfs), which is an important initial step in the process of modeling. Following, the application of Poisson- and Markov-type models for web traffic, inter-satellite routing, and model matching in the presence of long-range dependence is discussed. A time-discrete model provides insight into the impact of network stages on packet streams. Flow-level models, their parameters and applicability are discussed. The subsection on fluid models deals with the impact of modeling timescales on performance and the capability of detecting and describing the impact of bottlenecks. Finally, a couple of fractal-type models is presented.

2.4.1 PDF Estimation

In [87] the estimation of heavy-tailed probability density functions, their mixtures, and high quantiles is studied. First, the relevance of this issue in tele-traffic engineering is discussed and then a new combined estimation technique for such pdfs f is proposed. The “tail” of the pdf is estimated by a parametric tail model $f_{\gamma}^T(x) = \gamma x^{-\gamma-1} + 2\gamma x^{-2\gamma-1}$ and its “body” $f^B(t) = \frac{1}{x_{(k)}} \sum_{j=1}^N \lambda_j \varphi_j(\frac{t}{x_{(k)}})$ by a non-parametric method in terms of a finite linear combination of trigonometric functions $\varphi_k(t) = \sqrt{\frac{4}{\pi}} \cos((2k-1)\frac{\pi}{2}t)$ that are evaluated at the data points $x_{(k)}$ of a sample. To provide the minimum of the mean-squared error of the estimation, the parameters of the parametric and non-parametric parts are estimated by means of the bootstrap method and the structural risk minimization method. The latter parameters are determined by the number of extreme-valued data that are used in Hill’s estimate of the tail index γ and the number N of terms and coefficients of the linear combination. The new method is illustrated using some relevant mixtures of heavy-tailed pdfs, e.g., a mixture of a Gamma and a Pareto distribution arising from the delay modeling of IP traffic, and applied to construct a high quantile estimate. Furthermore, its effectiveness is shown by an application to real data derived from Web-traffic characteristics like session durations and volumes.

Measurements of Web traffic have shown that its characteristics like session durations, transferred volumes, and file sizes are governed by heavy-tailed distributions $F(x)$. In [88] the new task of Web data mining to estimate the un-

derlying heavy-tailed pdfs $f(x) = F'(x)$ by on-line algorithms is addressed. To guarantee a better estimation of the tail behavior, a new specific transformation scheme $T_{1/\gamma} : X \rightarrow [0, 1]$ is proposed. It is adapted to the empirical data $X = \{X_1, \dots, X_n\}$. Applying the latter, a new on-line estimator for such pdfs is developed as a further basic innovation. In this scheme the required extreme-value index $1/\gamma$ of the pdf f is reconstructed by a new recursive technique.

[31] tests empirical flow statistics against commonly used theoretical probability distributions such as hyper-exponential, Gamma, Weibull, Pareto and log-normal. It also discusses the suitability of statistical tests, such as the Kolmororov-Smirnov test and the Chi-Square test, for these kinds of tasks.

2.4.2 Poisson- and Markov-type models

Poisson cluster process for web traffic

In [89] is studied a model of HTTP sessions initiated at a Web server. The model is based on a previous analysis carried out by Liu et al. [90], among others, where one HTTP session is composed of one main object and possibly several in-line objects downloaded in parallel. The HTTP session finishes when the last downloading process ends. The objects of the study are the stationary distribution of the total number of objects which are currently downloaded and its covariance function. The model is a particular case of a Poisson cluster process and is more precisely defined as follows. At the Web server HTTP requests occur according to a Poisson process. To each HTTP request is associated a process of object requests determined by a transient Markovian Arrival Process, introduced in [91]. The Markovian assumption has the advantage of making the whole process very tractable and allows for much flexibility in modeling phenomena such as high correlations over long intervals of time [92], even though Markovian models do not exhibit long-range dependence *stricto sensu*. To each object is associated a mark, namely its downloading duration. We assume that the durations are iid random variables with a common distribution; this implies that the capacity of the system is sufficiently large that a transmission is not slowed down by other downloading processes which might be in progress at the same time.

Equivalent Poisson traffic

In the framework of inter-satellite routing, [93] presents an empirical traffic source model derived from an Internet backbone traffic trace, which is then compared to equivalent Poisson traffic as a point of reference. The comparison is done using traffic class dependent routing, which has the potential to differentiate between traffic classes using different optimization criteria in route calculation. Interestingly enough, the performance measures based on aggregate traffic flow show no significant difference between routing of empirical and equivalent Poisson traffic.

Markov Modulated Poisson Processes (MMPP)

Markov Modulated Poisson Processes (MMPPs) are very popular traffic models due to their ease of use (mainly due to their additive property) and to the availability of analytical results for the evaluation of their queueing behavior. In [94] a parameter fitting procedure using MMPPs that leads to accurate estimates of queueing behavior for network traffic exhibiting LRD behavior is proposed. The procedure matches both the autocovariance and marginal distribution of the counting process. A major feature is that the number of states is not fixed a priori, and can be adapted to the particular trace being modeled. The MMPP is constructed as a superposition of L 2-MMPPs and one M -MMPP. The 2-MMPPs are designed to match the autocovariance and the M -MMPP to match the marginal distribution. Each 2-MMPP models a specific time-scale of the data. The procedure starts by approximating the autocovariance by a weighed sum of exponential functions that model the autocovariance of the 2-MMPPs. The autocovariance tail can be adjusted to capture the long-range dependence characteristics of the traffic, up to the time-scales of interest to the system under study. The procedure then fits the M -MMPP parameters in order to match the marginal distribution, within the constraints imposed by the autocovariance matching. The number of states is also determined as part of this step. The final MMPP with $M \cdot 2^L$ states is obtained by superposing the L 2-MMPPs and the M -MMPP. The inference procedure was applied to traffic traces exhibiting long-range dependence, and its queueing behavior was evaluated through simulation. Very good results were obtained, both in terms of queueing behavior and number of states, for the traces used, which include the well-known Bellcore traces.

In [75] is proposed a simple MMPP traffic model that approximates the LRD characteristics of traffic traces measured at an edge router, at both flow and packet level, respectively. The MMPP model mimics the real behavior behind the interaction between users, protocols, and the network, using the notion of sessions and flows, therefore resulting in a simple and intuitive model. The estimated Hurst parameter of the synthetic traffic fits the one observed from data traces measured both at flow and packet level. The characteristics of the synthetic traffic generated with the model match the LRD characteristics observed in the measured traces over the time scales of interest. One of the interesting features of the proposed MMPP model is that it requires only five parameters. Three of these parameters can be directly mapped onto average traffic parameters, such as the average flow arrival rate, the average number of packets per flow, and the average arrival rate of packets within flows. The other two parameters define the notion of session, and are used to control the Hurst parameter of the synthetic traffic on the considered scaling range. The queuing behavior, for finite and infinite buffer queues, of the traffic generated by the model is coherent with the one of the measured traces at several different traffic loads.

While the model is not intended to offer an explanation of the reasons why Internet traffic is LRD, it does offer a simple and manageable tool for dimensioning and planning networks (link and buffer capacities), since the characteristics of the generated traffic are easily controlled through the model input parameters. The simplicity of the MMPP traffic model makes it an ideal traffic generator to drive simulations and, in addition, the Markovian properties of the MMPP traffic model make it analytically tractable, so that analytic solutions for simplified networking scenarios are possible.

2.4.3 Time-discrete models

In [95] the evolution of the interarrival and interdeparture times between voice packets when they are proceeding through a number of network nodes has been investigated. Instead of separately specifying the characteristics of each individual source, the arrival process in a node is modelled as the superposition of a single tagged stream and an independent background process that aggregates the remaining traffic sources. Because the load of a single voice stream is assumed to be very low compared to the load of the aggregate traffic, the tagged voice packets can be represented as markers (packets with size zero).

The tagged marker stream is characterized by the consecutive interarrival times between the markers, which are assumed to be identically distributed. The background arrival process on each network node is described on a slot-per-slot basis according to a general i.i.d. process. The results indicate that the mean interarrival time between markers remains constant throughout the network, and that both the variance of the interarrival time and the covariance between two successive interarrival times are increasing functions of the number of stages traversed by the markers. Also, it has been observed that if at the entrance of the first network stage two successive interarrival times are independent of each other, after one stage they become negatively correlated.

2.4.4 Flow-level models

Given the complexity of Internet traffic at packet level, some performance models are currently developed, particularly for elastic traffic under TCP control (see [40] and [38]), that consider higher level entities, namely the traffic “flows” or “sessions.” In order to investigate several assumptions commonly considered in such models, a detailed statistical analysis of the IP traffic processes at flow level is performed in [31], based on traces gathered on a wide area network. First, the concept of flow is thoroughly discussed and a method is suggested to estimate an adequate Time Out value to distinguish different flows within a traffic stream. Next, the main result of this contribution is the rejection of the Poisson hypothesis for the flow arrival process: it is shown that Gamma and Weibull distributions (the latter in agreement with [96]) provide excellent fits to the empirical flow inter-arrival time distributions; in addition, as can be seen on Figure 2.5, some consistent level of correlation is detected between successive flow arrivals. Finally, the study confirms the well-known heavy-tailed character (mostly Pareto-like) of flow size distributions, both for UDP and TCP traffic.

To evaluate the performance in scenarios of integration of streaming (voice, video) and elastic (data) traffic in a multiservice network, an analytical model is presented in [41]. Assuming Poisson flow arrivals on a network link and fair bandwidth sharing for elastic flows, the model is based on the performance of an M/G/1 Processor Sharing queue with time-varying capacity. The variable capacity is that which is left available by streaming flows, the packets of which are supposed to have priority with respect to those of elastic flows.

A related scenario is considered in [40] and [38], where two classes of




Figure 2.5: Autocorrelation functions of UDP (upper) and TCP (lower) flow inter-arrival time

elastic TCP traffic are distinguished, i.e., premium (high priority) and best effort; the TCP flows have a limited peak rate r . This scenario is modeled by a multiple server Processor Sharing queue with two priority classes. The analysis is presented in [40]. In particular, for the high priority traffic the mean sojourn time (i.e., file transfer time) is obtained from known analytical results for the M/G/C PS queue in [32]. For the best effort traffic, a simple accurate approximation for the mean sojourn time is derived. In [38] the analytical modeling results of [40] are compared with results obtained by ns simulations of the scenario described above.

With regards to TCP performance, PS models capture quite “roughly” the statistical multiplexing effect on flow level and allow for link dimensioning. However, these models do usually not provide insight into the impact of packet level parameters such as the round trip time and buffer size. Therefore, in [97] an *integrated* packet/flow level model for the analysis of a bottleneck link in a TCP/IP network environment is developed, combining the attractive features of both flow level and packet level models. The approach works roughly as follows. First, for a given number of flows, say n in the system, the throughput t_n is computed. This is done by using well known packet level TCP models reflecting the impact of round trip time and buffer size (see e.g. [98]). Next, on the flow level, the transfer times are obtained from a processor sharing model with Poisson flow arrivals and state dependent service rates t_n [32]. An interesting feature of this model is that the mean flow transfer time depends on the flow size distribution only through its mean value. This *insensitivity* property is confirmed by simulation.

[99] presents a similar integrated, hybrid packet/flow level modelling approach, but for the analysis of flow throughputs and transfer times in IEEE 802.11 WLANs. While the packet level model covers the statistics of the packet transfer at MAC level, the flow model reflects the system dynamics in terms of initiation and completion of flows. The resulting integrated model is analytically tractable and yields accurate approximations for throughput and flow transfer time values. In particular, a similar insensitivity property for the mean flow transfer time is obtained as in the TCP model of [97] discussed above.

2.4.5 Fluid models

Reference [100] investigates the impact of approximating packetized flows by fluid flows on the Complementary Cumulative Distributions (CCDs) of the unfinished work. For that purpose two discrete-time queueing models are compared. In the first model (referred to as the packet model) two timescales are present in the input traffic, i.e., a burst and a packet timescale. In the second one (referred to as the fluid model) only an (identical) burst timescale is present. It is assumed that time is divided in units of constant length, called frame times, and that every frame time is further divided into constant length time units, called packet times. As a typical model for a bursty traffic source a two-state discrete-time Markov source is considered. This Markov source

has a frame time as underlying time unit, i.e., the source can only change state on frame time boundaries. During a frame time in which the source is in the first, respectively second state, it generates a certain amount of bytes. In the packet model, these bytes are divided into fixed-length packets which are sent bunched up in the beginning of the frame time. In the fluid model the bytes are sent at a constant rate over the whole duration of the frame time. A queueing model with an infinite buffer is then considered in which the traffic of either M identical packet sources or M identical fluid sources with the same parameters is multiplexed. The performance measure of interest is the distribution of the amount of unfinished work in both systems. Numerical results show that the slope of the CCDs of the unfinished work obtained with both models is the same, i.e., fluctuations in the arrival pattern at the smallest timescale do not influence this slope. The probability that the amount of unfinished work is larger than a certain value is the smallest for the fluid model that neglects traffic fluctuations at the finest timescale. The difference between the two CCDs becomes however less important when the parameters of a scenario are changed such that the slope of the curves increases, i.e., when the probability of having a larger amount of unfinished work in the system becomes larger. Several properties of traffic scenarios that result in a larger slope of the curves are considered in [100]. In [101], the model presented in [100] is used in a case study dealing with a video streaming application.

The impact of time scales on the link capacity required for maintaining a desired QoS is discussed in [36], see also Figure 1.3.

In [102] the stochastic fluid flow model demonstrates the capability to reveal the impact of bottlenecks, i.e., shortages in capacity, on packet streams. This is achieved by comparing throughput histograms at the output with those at the input of the bottleneck. From these comparisons, it can be seen whether there is other interfering traffic sharing the bottleneck or whether the bottleneck has a buffer of significant size. This allows for both identification and classification of the bottleneck. It is furthermore observed that streams with different characteristics (constant or variable rate) inherit the same kinds of changes in their throughput histograms.

Based on [102], a measurement study [73] has been carried out on UDP packet streams carrying voice and video between Karlskrona, Sweden, and Würzburg, Germany. These videoconferencing streams passed a bottleneck, a 10 Mbps half-duplex link that also carried a disturbing UDP stream of vary-

ing intensity. As bottleneck indicator, throughput histogram difference plots were investigated. For the individual voice and video streams, these indicators displayed the behavior expected from [102]. Especially in case of the video stream, a “shared bottleneck” was already signalled at a level of disturbance when the user still did not feel any degradation in quality.

A fluid flow source model displaying multi-fractal properties is used in [103] and will be discussed further in section 2.4.6.

2.4.6 Fractal-type models

It has been observed in several measurements and statistical studies that network traffic exhibit fractal properties. After deep analyses it turned out that network traffic can also present multifractal characteristics. However, so far only a few relevant multifractal models have been developed and these are based on constructing a multiplicative process structure.

A fluid flow source model displaying multi-fractal properties was proposed by Mannersalo, Norros, and Riedi in the predecessor COST 257 Action [104] and is used in [103] to model data traffic. These properties are obtained by multiplying the outputs of a number N of Markov Modulated Rate Processes (MMRPs), each of which is acting on another timescale. Actually, one might think of one “fast” MMRP whose rate is modulated by “slower” MMRPs incorporating activities on different time scales. As N is finite, the whole process that is showing long-range dependence on short time scales becomes short-range dependent on long time scales. [103] presents time plots and variance-time log-log plots for different parameter settings.

Multifractal behavior was recently observed in several traces of IP WAN traffic. [105] proposes a novel traffic model that is able to capture multifractal behavior and characterizes jointly the processes of packet arrivals and packet sizes. The construction of the traffic process is based on stochastic L-Systems, introduced by biologist A. Lindenmayer as a method to model plant growth. They are characterized by an alphabet, an axiom, and a set of production rules. The alphabet is a set of symbols. The production rules define transformations of symbols into strings of symbols. Starting from an initial string (the axiom), an L-System iteratively constructs sequences of symbols, replacing each symbol by the corresponding string according to the production rules. The authors work with a single L-System alphabet and production rule, where the alphabet is a set of pairs, and each pair element represents a packet arrival rate and a

packet mean size. In this way, the traffic model is able to capture correlations between arrivals and sizes, leading to an accurate prediction of the queuing behavior. [105] provides a detailed comparison with a related multifractal model based on conservative cascades. The results, which include applying the fitting procedure to real observed data with multifractal scaling behavior on both the packet arrival and packet size processes, show that the L-System based model can achieve excellent fitting performance in terms of first and second order statistics and queuing behavior.

In [71] is presented the new monofractal model *Limit of the Integrated Superposition of Diffusion processes with Linear differential Generator* (LIS-DLG), which is not self-similar and has a more powerful modeling ability to capture scaling behavior when compared to self-similar models. It is argued that the monofractal model is flexible enough to accurately capture the fractal scaling of network traffic, and that there is no need to use more complex multifractal models. In this model the cumulants of the measured traffic are fitted to the cumulants of the process generated by the model, and it is shown that the resulting bispectrum of the traffic can be accurately captured. Several properties of the LISDLG model are presented in the report, including covariance structure, cumulants, spectrum, and bispectrum. The relevance and validation of the proposed model are demonstrated by application studies for measured Internet traffic.

2.5 Traffic Estimation

This final section focuses on the problems of inferring traffic on the basis of incomplete or indirect observations as well as on linear prediction based on recent measurements.

2.5.1 Traffic Inference

A Traffic Matrix (TM) reflects the volume of traffic that flows between source and destination nodes in a network. The nodes can refer to a variety of network elements such as POPs, routers, or even address prefixes. A POP-to-POP traffic matrix X captures the amount of traffic exchanged between two POPs, where X_{ij} represents the volume of traffic traveling from ingress POP i to egress POP j . The value of X_{ij} usually represents a bandwidth value

averaged over some time interval, although other types of elements are also possible. There are a number of traffic engineering tasks that could be greatly improved with the knowledge provided by traffic matrices. Capacity planning, routing protocol configuration, definition of load balancing policies, and fail-over strategies are tasks that would benefit from having information on the size and locality of traffic exchanges. An important example is the setting of OSPF or IS-IS routing weights. With knowledge of the TM, an algorithm for setting weights will select a routing that achieves a significantly better load balancing than one with an incorrect idea of the TM.

Obtaining a traffic matrix can be basically approached in two ways. One may directly measure it or one can rely on partial information to infer it. Measurement approaches have not been fully explored because they involve overcoming challenging engineering obstacles related to the deployment of a measurement infrastructure, and to the storage and processing of large amounts of information. Furthermore, the monetary cost may be high. Instead, previous work on obtaining traffic matrices has relied on statistical inference techniques that use partial information to estimate the TM. The term Network Tomography has been coined for this problem when the partial data come from repeated measurements of the traffic flowing along directed links in the network. Such data are usually obtained from SNMP, which allows measuring the total amount of incoming and outgoing bytes on a link typically over five-minute intervals. The idea behind inference approaches is to use these link statistics to infer the characteristics of end-to-end flows. End-to-end flows are defined within a single domain and are usually referred to as Origin-Destination (OD) pairs. In a POP-to-POP topology, the origin and destination nodes are POPs.

In addition to inference methods, it is also possible to formulate the traffic matrix estimation problem as a constrained optimization problem and use techniques such as Linear Programming (LP). In [106] is presented a comparative study of existing TM inference techniques. The evaluated statistical techniques are found to outperform an LP-based technique, still statistical techniques are significantly restricted in their ability to converge to the right solution. This is because they rely on scarce actual network information and they require intensive computation to reach reasonably accurate estimates. These restrictions impose a substantial burden on the quality of the starting point that should be provided to guide the estimation process. In [106] is used a very fast variant of the Expectation Maximization algorithm for the network tomography problem.

The improvements made are aimed at reducing the computation requirements of the algorithm, enabling it to expand the iterative horizon in search of global optima as solutions to the inference problem. Second, alternative modeling approaches are investigated to provide reasonable starting points for inference techniques. These starting points are called informed priors because they are obtained from models that incorporate substantial network information. The paper introduces choice models for generating informed priors. Different approaches are compared with respect to the estimation errors yielded, sensitivity to prior information required, and sensitivity to the statistical assumptions they make. The impact of characteristics such as path length and the amount of link sharing on the estimation errors is also studied.

[67] applies to cross traffic estimation a new methodology for analyzing and interpreting measurement collected over Internet. This new approach is based on inferring cross traffic characteristics that lead to the observed losses by using an associated a priori constructive model. The constructive model used in this paper is an MMPP/M/1/N single server bottleneck. The originality of this solution is that it starts with the observed loss process and infers inputs that have led to these observations. The methods presented in the paper provide a powerful solution to address the complexity of interpreting IP active measurement and empirical network modeling.

2.5.2 Traffic Prediction

In [107] is considered the usability of the traditional linear predictors especially in the case where the number of aggregated flows is moderate or small. First are stated some theoretical results based on the fractional Brownian motion, then real traffic traces with different aggregation levels are analyzed. Finally, some comments on the prediction based dynamical resource reservation are given. The main results can be summarized as follows. If the traffic has a power law-type variance structure and possibly non-stationary features, then the combined use of a moving average mean rate estimate and a fixed Fractional Brownian Motion (FBM) type variance function is motivated. This results in a robust and pretty good linear predictor. In the case of a prediction delay, a working engineering solution is to condition with respect to four or five geometrically increasing intervals with smallest interval about half the prediction interval. However, there are situations when the simple mean rate estimation is enough, i.e., it cannot be improved by more complicated algorithms.

Moreover, a straightforward application of traffic predictors in resource reservations may lead to problems as demonstrated in the case of ϵ -overallocation (see e.g. [108]).

[59] presents a fluid flow simulation experiment in which estimates of each connection's future load are used for online performance control of priority system. These estimates are fed to a rate allocation algorithm that decides on how the available rate will be shared among low-priority queues. The objective of this algorithm is to minimize packet loss by keeping all low-priority queues at the same size. The estimates are delivered by an adaptive linear filter, acting on time slots of 100 ms. Also, the influence of the number of coefficients of the filter, i.e., window size or time horizon on the obtained predictions is discussed.

Chapter 3

Queueing Models

Sabine Wittevrongel
Ghent University, Belgium

Contributors:

Nail Akar (Bilkent University, Turkey), Carlos Belo (Telecommunications Institute, Portugal), Ana da Silva Soares (Université Libre de Bruxelles, Belgium), Markus Fiedler (Blekinge Institute of Technology, Sweden), Dieter Fiems (Ghent University, Belgium), Peixia Gao (Ghent University, Belgium), Veronique Inghelbrecht (Ghent University, Belgium), Udo Krieger (Otto Friedrich University Bamberg, Germany), Koenraad Laevens (Ghent University, Belgium), Michel Mandjes (Center for Mathematics and Computer Science, Netherlands), Ilkka Norros (VTT Information Technology, Finland), Olav Østerbø (Telenor Research & Development, Norway), Detlef Sass (University of Stuttgart, Germany), Kathleen Spaey (University of Antwerp, Belgium), Hung Tran (Telecommunications Research Center Vienna, Austria), Hans van den Berg (TNO Telecom, Netherlands), Jorma Virtamo (Helsinki University of Technology, Finland), Joris Walraevens (Ghent University, Belgium), Sabine Wittevrongel (Ghent University, Belgium)

3.1 Introduction

In this chapter is presented a brief overview of the queueing models that have been studied in the COST279 project. Specifically, in Section 3.2 a number of specific discrete-time queueing models is discussed. Section 3.3 addresses some new developments with respect to fluid flow models. Section 3.4 summarizes the work on Gaussian storages. In Section 3.5 some new results on processor sharing models are presented. Section 3.6 is devoted to the analysis of some other continuous-time queueing models. Finally, in Section 3.7 some techniques to study the behavior of a network of queues are briefly discussed.

3.2 Discrete-time queueing models

In a discrete-time queueing model the time axis is assumed to be divided into fixed-length intervals, usually referred to as slots. This section gives an overview of a number of specific discrete-time queueing models studied in the COST-279 project, as well as the results obtained for these models. For the analysis of the models, analytical techniques mainly based on extensive use of probability generating functions (pgfs) have been developed. For more details on the developed analysis techniques the interested reader is referred to the corresponding Temporary Documents.

Queues with priority scheduling

Priority scheduling is a hot topic in multimedia networks. For real-time applications, it is important that the mean delay and the delay jitter are not too large. Therefore, this type of traffic is given priority over non-real-time traffic in the switches/routers of the network, i.e., delay-insensitive traffic is serviced in a switch only when no delay-sensitive traffic is present.

In [109] is considered a discrete-time single-server queueing system with infinite buffer space, a head-of-line (HOL) priority scheduling discipline, and a general number of priority classes. All types of cell arrivals are assumed to be independent and identically distributed (i.i.d.) from slot-to-slot, but within one slot the number of cell arrivals from different classes can be correlated. The system has one server that provides the transmission of cells at a rate of 1 cell per slot, i.e., the service times of the cells are deterministically equal to one slot (which is a sufficient assumption for ATM networks). First, an expression for the joint pgf of the system contents of all priority classes is derived. From this joint pgf, the marginal pgfs of the system contents of each priority class separately and of the total system contents are found. Furthermore, the pgf of the cell delay of each class is calculated. The analysis of the latter is largely based on the concept of sub-busy periods. From the generating functions obtained, performance measures (the moments and approximate tail probabilities of system contents and cell delay) are derived. Especially the analysis of the tail behavior is an important result of [109]. It is shown that the tail behavior is not necessarily exponential. Numerical results illustrate the impact and significance of priority scheduling in an ATM switch. For two priority classes, this analysis has also been extended to the case of general packet service times.

Specifically, in [110, 111] general packet service times and a preemptive priority discipline are considered, whereas [112] considers general service times and non-preemptive priorities.

Another way to study queues with priority scheduling is discussed in [113]. Specifically, in [113] a discrete-time single-server queue subjected to server interruptions is investigated. Server interruptions are modelled as an on/off process with geometrically distributed on-periods and generally distributed off-periods. The messages that arrive in the system possibly require more than one slot of service time, implying that a service interruption can occur while a message is in service. Therefore, different operation modes are considered, depending on whether service of an interrupted message continues, partially restarts, or completely restarts after an interruption. For each alternative, expressions for the steady-state pgfs of the buffer contents at message departure times and at random slot boundaries, of the unfinished work at random slot boundaries, of the message delay, and of the lengths of the idle and busy periods are established. From these results, closed-form expressions for various performance measures, such as mean and variance of the buffer occupancy and the message delay, are derived. Numerical results show the deterioration of system performance caused by service repetitions. In particular, it is observed that the mean length of the available-periods crucially determines the system stability for the (partial) repeat after interruption mode.

The model considered in [113] can be used to assess the performance of a multi-class preemptive priority scheduling system. In this case, the system interrupts service of lower class messages to serve higher class messages. Assume that class i has a higher priority than class j , for $i < j$. Then, class 2 messages receive service during the idle periods of class 1 messages. Class 3 messages are served during the idle periods of class 2 messages (the busy periods include the interruptions) and so on. The continue after interruption mode and the repeat after interruption mode then correspond to preemptive resume and preemptive repeat priority scheduling. The partial repeat after interruption operation mode can be considered as an intermediate case between both types of priority scheduling and allows to investigate the influence of packetizing in preemptive priority systems.

Queues with server vacations

Queueing systems with server vacations have proven to be a useful abstraction of systems where several classes of customers share a common resource, such as priority systems and polling systems, or of systems where this resource is unreliable, such as maintenance models and Automatic Repeat Request (ARQ) systems.

A discrete-time gated vacation system is considered in [114]. The classical gated vacation system can be seen as one with two queues separated by a gate. Arrivals are routed to the queue before the gate whereas customers in the queue after the gate are served on a first-in-first-out basis. When there are no more customers in the latter queue, the server takes a vacation and opens the gate – all customers move instantaneously from the queue before to the queue after the gate – upon returning from this vacation. In [114], the classical gated vacation queueing system is extended in the sense that it is also allowed that customer arrivals are immediately routed to the queue after the gate. The model under investigation allows to capture performance of, among others, the exhaustive and the gated queueing systems with multiple or single vacations. Using a generating-functions approach and the method of the supplementary variable, expressions are obtained for performance measures such as the moments of the system contents at various epochs in equilibrium and of the customer delay. The results depend on a constant that has to be determined numerically.

Multiserver queues with geometric service times

In most of the existing literature on discrete-time multiserver queueing models, the service times of customers are assumed to be constant, equal to one slot or multiple slots. In [115], a discrete-time multiserver queue with geometric service times, an infinite buffer size, a first-come-first-served (FCFS) service discipline, and general independent packet arrivals is considered. The behavior of the queueing system is studied analytically by means of a generating-functions approach. This results in closed-form expressions for the pgfs of the system contents and the packet delay. Furthermore, these pgfs are used to derive various additional performance measures. Specifically, expressions are derived for the mean values, the variances, and the tail probabilities of the system contents and the delay. In [116] the analysis is further extended from the case of an

uncorrelated packet arrival process to the case of a two-state correlated arrival process. The delay analysis is based on a general relationship between system contents and packet delay established in [117], valid for any discrete-time multiserver system with geometric service times, regardless of the exact nature of the arrival process.

Queues with bursty traffic

In [100] the authors compare two discrete-time queueing models: a *packet model*, where two timescales are present in the input traffic, i.e., a burst and a packet timescale, and a *fluid model*, where only an (identical) burst timescale is present. Time is divided in units of constant length, called frame times, and every frame time is further divided into constant length time units, called packet times. As a typical model for a bursty traffic source, a two-state discrete-time Markov source is considered. This Markov source has a frame time as underlying time unit, i.e., the source can only change state on frame time boundaries. During a frame time in which the source is in the first, respectively second state, it generates a certain amount of bytes. In the packet model, these bytes are divided into fixed-length packets that are sent bunched up in the beginning of the frame time. In the fluid model the bytes are sent at a constant rate over the whole duration of the frame time. A queueing model with an infinite buffer is then considered in which the traffic of either M identical packet sources or M identical fluid sources with the same parameters is multiplexed. The performance measure of interest is the distribution of the amount of unfinished work in both systems.

Of special interest in [100] is the impact of approximating packetized flows by fluid flows on the CCDs of the unfinished work. Numerical results show that the slope of the CCDs of the unfinished work obtained with both models is the same, i.e., fluctuations in the arrival pattern at the smallest timescale do not influence this slope. The probability that the amount of unfinished work is larger than a certain value is the smallest for the fluid model, which neglects traffic fluctuations at the finest timescale. The difference between the two CCDs becomes however less important when the parameters of a scenario are changed such that the slope of the curves increases, i.e., when the probability of having a larger amount of unfinished work in the system becomes larger.

Models for optical buffers

Optical packet switching and optical burst switching seem promising techniques to cope with the explosive growth of the Internet traffic. In the design of all-optical switches, the lack of optical Random Access Memory (RAM) poses a big challenge. Besides wavelength conversion and deflection routing, the use of fiber delay lines (FDLs) can help alleviate the output port contention problem. These FDLs are passive components that can delay an optical packet or an optical data burst for a fixed time. Usually, an FDL buffer implements the delays $0 \cdot D, 1 \cdot D, \dots, M \cdot D$, where D is the so-called granularity and $M \cdot D$ can be considered as the capacity of the FDL buffer. Note that thus not all delays can be obtained, typically leading to the creation of voids in the scheduling and to an underutilization of the output channel. For this reason, FDL buffers are also sometimes called degenerate buffers.

The performance of an FDL buffer is analyzed in [118]. The quantity of interest in the analysis is the so-called scheduling horizon. It is defined as the earliest time at which the channel will become available again, and can be considered the equivalent of the unfinished work in non-degenerate buffers. If one denotes by H_k this scheduling horizon as seen by the k -th arrival, one can easily establish—assuming an infinite FDL buffer—the following recursion :

$$H_{k+1} = \left[B_k + D \left\lceil \frac{H_k}{D} \right\rceil - \tau_k \right]^+. \quad (3.1)$$

Here B_k denotes the size of the k -th burst, and τ_k the interarrival time between the k -th and $(k+1)$ -th burst. One easily recognizes part of the evolution equation for non-degenerate buffers, involving the operation

$$[\dots - \tau_k]^+. \quad (3.2)$$

Under the usual assumptions of i.i.d. interarrival times and i.i.d. burst sizes, the solution to this problem in the transform domain is well-known. The part

$$D \left\lceil \frac{H_k}{D} \right\rceil \quad (3.3)$$

reflects the finite granularity of the FDLs. By using an identity involving the complex D -th roots of unity, this operation on random variables can be translated into an operation on their pgfs. By combining both partial solutions, one

obtains in the end the pgf of the scheduling horizon H in equilibrium. From the latter, one can obtain several measures of interest, such as the maximum tolerable load, i.e., the load at which the infinite system becomes unstable. Due to the creation of voids, this load is typically less than unity; it also shows a slight dependency on the burst size distribution. Further, a heuristic can be used to map the tail probabilities $\text{Pr}[H > M \cdot D]$ to loss probabilities in a finite system of capacity $M \cdot D$. An optimum granularity D_{opt} exists, establishing a compromise between increasing capacity ($D \rightarrow \infty$) and small voids ($D \rightarrow 0$). This optimal value not only depends on the burst size distribution, but also on the load offered to the system.

In [119] the performance of an optical packet switch is investigated. Specifically, an FDL-structure consisting of N delay lines with increasing lengths is considered. It is assumed that the optical packets have deterministic lengths and the i -th delay line ($i = 1, \dots, N$) has a length of i times the packet length. Time is slotted, where one slot corresponds to the time needed to transmit a packet. The number of packet arrivals are assumed to be i.i.d. from slot-to-slot. The scheduling discipline is smallest FDL first. In this scheduling discipline, if i packets arrive at the same time, they are put in the i delay lines with the smallest lengths (if $i > N$, $i - N$ packets are lost). First, an expression for the pgf of the steady-state delay line contents of the FDL-structure with increasing lengths is derived. From this pgf, the packet loss rate (PLR) in an output buffering optical packet switch is calculated. Through some figures, the impact of the number of delay lines and the load on the PLR is shown. An important conclusion is that putting one or two delay lines at each output can reduce the PLR significantly. Adding even more delay lines per output does not reduce the PLR significantly though. If lower PLR's have to be obtained, the scheduling discipline discussed in [119] (smallest delay line first) is not sufficient and more complex scheduling methods are necessary.

Burstification queues

In the edge routers of an OBS network, IP packets are assembled into bursts. Core OBS routers forward these bursts in the optical domain through the OBS network. An OBS edge router can be decomposed into multiple burstification units (BU). Each BU consists of a set of separate output queues. In [120], the burstification of a single isolated output burst queue is investigated. First, a single-threshold burst assembly mechanism is studied. Bursts will be released

if they contain exactly S packets. In this case, especially for low throughputs, the packets may have long delays. Therefore, as a next step, also a two-threshold model is investigated, where besides a threshold on size, a threshold on a burst's age is imposed. Thus, bursts will also release if since the start of their assembly a time T has expired. For both queueing models, some performance characteristics are calculated. Results include the probability mass function (pmf) of the system contents in the output queue of the OBS edge routers, the pmf of the delay of the bursts (defined as the interdeparture time of two bursts), and the pmf of the delay of the individual packets. Using these pmfs, the mean values, the variances and the tail distributions of the system contents, the burst and packet delay are derived.

Transmitter queue of a stop-and-wait ARQ system

In [121] an analytical approach for studying the queue length and the packet delay in the transmitter buffer of a system working under a stop-and-wait retransmission protocol is presented. The buffer at the transmitter side is modelled as a discrete-time queue with an infinite storage capacity. The numbers of packets entering the buffer during consecutive slots are assumed to be i.i.d. random variables. The information packets are sent through an unreliable and non-stationary channel, modelled by means of a two-state Markov chain. An explicit formula is derived for the pgf of the system contents. This pgf is then used to derive several queue-length characteristics as well as the mean packet delay.

3.3 Fluid flow models

In a fluid flow model (FFM) the amount of work delivered to a queue or processed by a server is modelled as a continuous-time flow. A fluid queue is generally solved by first finding the eigenvalues and eigenvectors of the underlying differential system and then obtaining the coefficients of the associated spectral expansion by solving a linear matrix equation. This section presents some new COST279 developments with respect to fluid flow analysis.

Algorithmic approach

Consider a Markov modulated fluid queue, i.e., a two-dimensional continuous-time Markov process $\{(X(t), \varphi(t)) : t \in \mathbb{R}^+\}$ where $X(t)$ takes values in \mathbb{R}^+ and $\varphi(t)$ in \mathcal{S} , a finite set. The component $X(\cdot)$ is called the *level* and $\varphi(\cdot)$ is called the *phase*. The level is subordinated to the phase in the following way. The phase process $\{\varphi(t) : t \in \mathbb{R}^+\}$ is an irreducible Markovian process. During those intervals of time in which the phase is constant, say equal to i , the level increases or decreases at a constant rate dependent on i , or it remains constant. If $X(t) = 0$ and if the rate at time t is negative, then the level remains at 0.

Ramaswami [122] shows that $\{X(t)\}$ has a phase-type stationary distribution using the dual process of $\{(X(t), \varphi(t))\}$. Most importantly, he also constructs a very efficient computational procedure, based on the logarithmic-reduction algorithm of Latouche and Ramaswami [123] for discrete-level Quasi-Birth-Death (QBD) processes: he thereby reduces a complex continuous time, continuous state space problem to a familiar, simple discrete time, discrete state space system. The use of the dual process in [122] is motivated by a property that relates the stationary distribution of the original $\{(X(t), \varphi(t))\}$ process to first passage probabilities at level 0 in the dual process. In [124] the similarities with QBDs are reinforced by showing that one may actually *directly* use these first passage probabilities in the original process. Also, another probabilistic interpretation of Ramaswami's computational procedure is given.

Large-scale finite fluid queues

Except for some structured models, e.g., the Anick-Mitra-Sondhi (AMS) fluid flow model [125], it is in general hard to find the eigensystem in a computationally efficient and stable manner. Moreover, the linear matrix equation to solve for the coefficients in the finite fluid queue case is known to be ill-conditioned, especially in the case of large buffer sizes.

In [126], a numerically efficient and stable method for solving large-scale finite Markov fluid queues is developed. No special structure is imposed on the underlying continuous-time Markov chain, i.e., the eigenvalues and eigenvectors need not be determined in closed form. The authors propose an alternative method that relies on decomposing the differential system into its stable

forward) and anti-stable (backward) subsystems (as opposed to finding eigenvalues) using a method that is based on the additive decomposition of a matrix pair with respect to the imaginary axis. There are a variety of numerical linear algebra techniques (with publicly available codes) that can be used for such an additive decomposition, including the generalized Newton iterations, the generalized Schur decomposition, and the spectral divide and conquer methods. Using the generalized Newton iterations, which have quadratic convergence rates, it is shown in [126] that the accuracy of the proposed method does not depend on the buffer sizes and that in the limit the finite fluid queue solution converges to that of the infinite fluid queue. Moreover, it is demonstrated that fluid queues with thousands of states are efficiently solvable using the proposed algorithm.

Voice and multi-fractal data traffic

The FFM has shown being able to incorporate many types of traffic, i.e., superpose them for analysis in a unified model. In [103] it is augmented by a model displaying multi-fractal behavior, which is described in Section 2.4.6. This model can be matched to the characteristics of real traffic by choosing the appropriate parameters for the sub-processes. [103] investigates the interaction between multi-fractal data traffic, and voice traffic with suppressed silence phases and consideration of the on-hook-state of the Internet phone, modelled by a 3-state on-off model. The fluid flow calculations can be used in a straight-forward manner, with the only exception that the pseudo-rates of the sub-processes contributing to the data traffic process are multiplied instead of added. As a result, formulas expressing queuing delay and loss as experienced individually by voice and data are obtained. A case study has been carried out, investigating the maximal load under given delay quantiles. As expected, this load level depends heavily on the variance of the data traffic. In general, the voice traffic yields better performance in terms of loss and delay and helps to increase the maximal load while still meeting the target performance values.

Superposition of general ON-OFF sources

In [127], the effect of a superposition of general ON-OFF sources on a multiplexer is studied. Sources are statistically identical and independent. During the ON period a source emits at a constant rate either in a fluid-flow fashion or

by periodically emitting fixed size packets. During the OFF period the source remains silent. Both the ON and OFF periods are random variables with general distributions but finite mean values. The distributions considered include distributions of the Pareto type, which are known to lead to traffic having the LRD property.

The superposition of a number of such general ON-OFF sources results in a stochastic process called semi-birth and death (semi-BD). The state of the semi-BD is the number of active sources at a given time; the random variable of interest is the holding time in that state. In this case, the traffic generated simply equals the number of active sources times the individual rate. For the case of the semi-BD, the stationary distribution of the holding time is given in [127] in explicit form as a function of the state considered, the number of sources, and the distribution of the ON and OFF periods. It is further argued that, for the semi-BD, the distribution of the holding time would tend to an exponential even with a moderate number of sources. The above result strongly suggests that the now classical AMS solution for the probability of buffer overflow of an infinite buffer multiplexer with a superposition of exponential ON-OFF sources as input could be applied to the case of general ON-OFF distributions. The remaining of [127] is devoted to evaluating, theoretically and by way of simulation, the circumstances under which the AMS solution is a good approximation as a function of the number of sources, the distributions of the ON and OFF periods, and the desired level of probability of overflow. For more details, the reader is referred to [127].

Feedback fluid queues

In [128] is considered a single point in an access network where several bursty users are multiplexed. The users adapt their sending rates based on feedback from the access multiplexer. Important parameters are the user's peak transmission rate p , which is the access line speed, the user's guaranteed minimum rate r , and the bound ϵ on the fraction of lost data. Two feedback schemes are proposed and studied. In both schemes the users are allowed to send at rate p if the system is relatively lightly loaded, at rate r during periods of congestion, and at a rate between r and p , in an intermediate region. For both feedback schemes an exact analysis is presented, under the assumption that the users' file sizes and think times have exponential distributions. The techniques are used to design the schemes jointly with admission control,

i.e., the selection of the number of admissible users, to maximize throughput for given p , r , and ϵ . Next is considered the case where the number of users is large. Under a specific (many-sources) scaling, explicit large deviations asymptotics for both models are derived. The extension to general distributions of user data and think times is discussed.

The model is also extended to a “buffer-dependent” Markov fluid queue, defined as follows. A Markov fluid source is defined as a (continuous-time) transition matrix Q of dimension d , and a traffic rate vector r (describing the generation of traffic at a constant fluid rate $r(i)$ when the Markov chain is in state i , for $i = 1, \dots, d$). Now consider N of these Markov fluid sources of dimension d , i.e., (Q_n, r_n) , for $n = 1, \dots, N$, and suppose traffic is generated according to the n th Markov fluid source when the buffer level is between B_{n-1} and B_n (where the B_n are increasing, and $B_0 = 0$; B_N can be ∞). For this model, the complete buffer contents distribution is derived in terms of the solution of an eigensystem.

Fair queueing systems

In [129], an FFM for a fair queueing system with unidirectional coupling for several different classes is considered, where each class has a predefined minimum bandwidth guaranteed. These minimum bandwidths form a decomposition of the total bandwidth and avoid starvation of a class caused by another class. A multiplexing gain is achieved by passing down the residual bandwidth of a class to the lower adjacent class. Therefore, this system is called *unidirectionally coupled*. Also, each class has its own FIFO buffer for exclusive usage. The system is described by an FFM, and hence sources and server are assumed to be Markov modulated fluid processes (MMFP). The key observation concerning the residual bandwidth is that a class, lending its residual bandwidth, is oblivious to this. Also receiving residual bandwidth is, from the receiving class point of view, just an additional server process. This additional server process is interpreted to be, again, an MMFP, and certain states and the buffer’s mean busy period of the lending adjacent class are used to model this additional process. The states mentioned are the under-load states, because state transition among these states reflects—assuming the buffer is empty—the dynamics of the residual bandwidth. The transition matrix of the additional server process

can be built with it. The distribution of the buffer content is found by applying the FFM. To determine the distribution for a specific buffer the distribution of the previous adjacent buffer has to be calculated, apart from the first buffer. Also, two estimates for the overflow probability of a buffer are obtained. First, a more accurate estimate is calculated by applying the full FFM, i.e., all eigenvalues and eigenvectors are used. Second, a fast and robust estimate is found by using only the dominant eigenvalue and the Chernov large deviation bound. This is an conservative estimate with larger deviation than the other estimate.

Bottleneck identification and classification

The stochastic FFM [125] has also shown to be capable of revealing the impact of bottlenecks, i.e., shortages in capacity, on packet streams by comparing bit rate histograms at the output with those at the input of the bottleneck [102]. From standard fluid flow analysis for MMPPs, the joint probabilities that the buffer is empty or non-empty in each state are calculated. As shown in [102], these probabilities are the key for obtaining the output bit rate distribution both for individual and total traffic streams. From comparisons with input bit rate distributions, it can be seen whether there is interfering traffic in the bottleneck or whether the bottleneck has a buffer of significant size. This allows not only for an identification, but also for a classification of the bottleneck. While the maximal capacity of the bottleneck is revealed in the output bit rate histogram of the total stream, it may under certain conditions also become visible in the corresponding histogram of an individual stream.

3.4 Gaussian storages

This section overviews some new developments with respect to the most probable path technique to derive estimates of the queueing performance for queues with Gaussian input traffic. Also, a method to determine delay quantiles of a multiplexer with Gaussian input—involving a fitting procedure to Ornstein-Uhlenbeck processes—is discussed.

Most probable path technique

By the theory of large deviations of Gaussian processes, the probability that a simple queue with Gaussian input exceeds a level x can be approximated by

$$\Pr[Q \geq x] \approx \exp\left(-\frac{1}{2}\|f_x\|_R^2\right), \quad (3.4)$$

where f_x is the path of the input process that creates a queue of size x at time 0 and has the smallest reproducing kernel Hilbert space (RKHS) norm $\|\cdot\|_R^2$ among all paths creating such a queue.

The framework was generalized in a straightforward way to a two-class priority system already in the preceding COST 257 Action as follows. Assume that the two arrival processes are independent continuous Gaussian processes with stationary increments. Consider the most probable *path pair* that creates a total queue (both classes together) of size x . If this path of the higher priority input does not create a positive queue at time 0, then this path pair is also the most probable one to create a lower class queue of size x .

The remaining case is studied in [130]. The idea is to compute an easily characterized heuristic approximation to the most probable path pair, where the higher priority traffic essentially fills the link while the lower class traffic is accumulating in the queue. The same principles can be applied to a generalized processor sharing system with two classes by replacing link filling by filling the quota guaranteed for a traffic class. Simulations showed that these approximations were sufficiently accurate for many practical purposes, like for studying the effects of setting GPS weights.

In [108], another kind of modification of the basic Gaussian queue is studied. The service capacity is not any more constant, but continuously varied according to the observed traffic rate, with a constant delay Δ . The allocated capacity is $1 + \epsilon$ times the observed rate. That is, the cumulative service process is defined as

$$C_t = (1 + \epsilon)(A_{t-\Delta} - A_{-\Delta}), \quad (3.5)$$

where A is the cumulative input process. The queue length process is defined as a supremum of the net input process:

$$Q_t = \sup_{s \leq t} ((A_t - A_s) - (C_t - C_s)). \quad (3.6)$$

Since the net input process is Gaussian, the basic estimates of the queue length distribution and the most probable paths are directly available.

Figure 3.1 shows, in case of fractional Brownian input, the most probable path that creates a queue of size 4 at time 0. Note how the input process “fools” the prediction by making the input first very slow and then, when the control cannot react any more, suddenly speeding up.




Figure 3.1: The input rate of the most probable way to obtain a queue of size 4, when service capacity is varying according to traffic prediction. The input process is a fBm with Hurst parameter 0.75.

The paper [131] makes a substantial new contribution to the most probable path approach described above by establishing a technique for identifying most probable paths that are truly infinite-dimensional combinations of covariance functions. This paper focuses on the simplest of this kind of problems. Consider a simple Gaussian queue with centered input process Z and service rate 1. What path β^* is the most probable one among input paths f that produce a busy period starting at 0 and straddling the interval $[0, 1]$, that is, satisfy $f(t) \geq t$ for all $t \in [0, 1]$?

The crucial observation is that there exists a closed set $S^* \subseteq [0, 1]$ such that (under certain conditions that were erroneously overlooked in the TD; the

result holds for fBm, and non-smoothness of the paths of Z may suffice) $\beta^*(t) = \mathbb{E}[Z_t | Z_s = s \forall s \in S^*]$. Thus, the most probable path is determined by those of its points where the criterion of being at or above the diagonal is met sharply. For Brownian motion, it is well known that $S^* = \{1\}$. In the case of fBm, it turns out that S^* has the form $[0, s^*] \cup \{1\}$ (resp. $[s^*, 1]$) if the self-similarity parameter H is larger (resp. smaller) than $1/2$.

Delay quantiles

In [132], delay quantiles are derived for the Gaussian voice traffic model. Delay quantiles γ_k with $\Pr[\text{delay} > \gamma_k] = 10^{-k}$ play an important role when it comes to dimensioning and performance evaluation. Especially large links should be dimensioned such that delay quantiles remain relatively small. However, as related methods often suffer from time and memory limitations together with numerical instabilities, it is not necessarily simple to obtain numerical results for systems incorporating many streams. The superposition of a large number of homogeneous Markovian on-off sources asymptotically approaches an Ornstein-Uhlenbeck process (OUP) representing a Gaussian process with exponential autocorrelation function. For such an OUP/D/1 model, [132] derives a closed-form expression for delay quantiles as

$$\gamma_k \simeq \frac{\sigma_R}{\omega_R C m} \left(k \log(10) - \frac{7}{4} m - \left(\frac{k}{30} + \frac{1}{4} \right) m^2 \right), \quad (3.7)$$

where $m = \frac{C - \mu_R}{\sigma_R}$. The parameter μ_R represents the mean rate, σ_R is the corresponding standard deviation, ω_R denotes the reciprocal time constant of the autocorrelation function of the rate, and C stands for the capacity of the link. By comparing with—in sense of the model—exact results for a non-finite number of fluid flow on-off sources, we find that for the relevant parameter region $\Pr[\text{source on}] \simeq 0.4$, $k \in [3, 6]$, $m \in [0, 1.6]$ and $N > 100$, the related deviations of the approximated delay quantiles do not exceed 10 %, which makes (3.7) a well-working approximation formula.

3.5 Processor sharing models

PS models are widely applicable to situations in which different users receive a share of a scarce common system resource. In particular, over the

past few decades PS models have found many applications in the field of the performance evaluation of computer-communication systems. The standard PS model consists of a single server assigning each customer a fraction $1/n$ of the service rate when there are n customers in the system. Cohen [32] generalizes the PS-model to the so-called GPS model, where each customer receives a fraction $f(n)$ of the service speed when there are n customers at a node, where $f(\cdot)$ is an arbitrary function (under weak assumptions). The GPS model significantly enhances the modeling capabilities of the PS model. Interestingly, over the past few years the GPS model studied by Cohen in [32] in 1979 has received a renewed interest in the literature on performance of computer-communication networks (see, e.g., [133], [134], [135]). A particularly attractive feature of (G)PS models is that in many applications they cover the main factors determining performance, and on the other hand, are still simple enough to be analytically tractable (see, e.g., the analysis in [136], [137], [32]).

Sojourn times for PS models with multiple servers and priority queueing

In [40] is studied a PS model with multiple servers and two priority classes (without loss of generality). When the number of high-priority customers does not exceed C (the number of servers), each high priority customer occupies a single server and is served at unit rate. When the number of high-priority customers is larger than C , the system switches to a processor sharing mode and the total service capacity C is equally shared among the high priority customers. The service process of low priority customers proceeds in a similar way, but with two specific restrictions: (1) high priority customers have strict priority (in a preemptive-resume fashion) over low priority customers, and (2) at any moment in time the only servers available to low priority customers are those servers that are not used by high priority customers at that moment.

The motivation for the study of this model is its application in packet switched networks supporting service differentiation by serving the packets of high quality flows with strict priority over packets of low quality (“best effort”) flows in the network nodes (e.g., routers in an IP network). Data traffic flows are usually subjected to a network flow control mechanism such as TCP. When bandwidth is temporarily limited, the flow control mechanism decreases

the transmission rate of each of the elastic flows, assigning each of them a fair share of the available bandwidth. The bandwidth effectively consumed by each flow may be limited by, e.g., the access line speed (modem speed), which is modelled by the service speed of a single server in the model described above. The sojourn times in the model represent the download times of elastic TCP flows, e.g., resulting from the download of a file or a web page.

In [40] the mean sojourn times in the multiserver queueing model with PS service discipline and two priority classes described above are studied. For the high priority class, closed-form expressions for the mean sojourn times are presented in a general parameter setting, based on known results for the GPS model (see [32]). For low priority customers, closed-form expressions are derived for several special cases: the single-server case where the service times of the low-priority customers are exponentially distributed, and the multiple-server case with exponential service times with the same means. In all other cases, exact explicit expressions for the mean sojourn times of the low priority customers cannot be obtained. Therefore, a simple and explicit approximation is proposed and tested. Numerical results demonstrate that the approximation is accurate for a broad range of parameter settings. As a by-product, it is observed that the mean sojourn times of the low-priority customers tend to decrease when the variability of the service times of the low-priority customers increases.

Application of these results to the setting with elastic TCP traffic is given in [38].

Throughput measures for PS models

In [138] the authors have specified, derived, and compared a set of throughput measures in PS queueing systems modeling a network link carrying elastic TCP data calls (e.g., file downloads). The available service capacity is either fixed, corresponding to a stand-alone dedicated GPRS network, or randomly varying, corresponding to an integrated services network where the elastic calls utilize the capacity left idle by prioritized stream (e.g., speech) traffic.

While from the customer's perspective the *call-average throughput* is the most relevant throughput measure, in PS systems the call-average throughput may be hard to determine analytically, and this is an important reason to assess the closeness of a number of other throughput measures. In several papers the *time-average throughput* ([139], [140], [141]), defined as the ex-

pected throughput the “server” provides to an elastic call at an arbitrary (non-idle) time instant, or the *ratio* of the expected transfer volume and the expected sojourn time ([142], [143], [144], [145]) are applied to approximate the call-average throughput.

In [138] is introduced the *expected instantaneous throughput*, i.e., the throughput an admitted call experiences immediately upon admission to the system, as a new throughput measure, which can be analysed relatively easily. The experiments demonstrate that the newly proposed expected instantaneous throughput measure is the *only* one among these throughput measures that excellently approximates the call-average throughput for each of the investigated PS models and over the entire range of elastic traffic loads. In particular for the model integrating speech and data traffic the other throughput measures, such as the time-average throughput or the ratio of the expected call size and the expected sojourn time, significantly underestimate the call-average throughput. An intuitive reasoning for the generally (near-)perfect fit of the expected instantaneous throughput is that, apparently, the throughput an elastic call experiences immediately upon arrival is an excellent predictor of what the call is likely to experience throughout its lifetime. Moreover, among the considered throughput measures, the expected instantaneous throughput is the *only* approximate measure that is truly *call-centric*.

The numerical experiments further reveal that the expected call-average throughput of elastic calls in the considered PS models is to a considerable degree *insensitive* to both the variability of the available capacity and the call duration distribution. This insensitivity does not hold if the data performance is measured by the expected sojourn time.

3.6 Other continuous-time queueing models

This section discusses a number of continuous-time queueing models and associated analysis techniques, studied in the COST-279 project.

Instantaneous and averaged queue length in an M/M/1/K queue

In [146] the authors study the dynamics of the joint process of the instantaneous queue length $L(t)$ of an $M/M/1/K$ system together with the expo-

nentially averaged queue length $S(t) = \int_0^\infty L(t-u)\alpha e^{-\alpha u} du$, where α is a weighing parameter. The arrival and service rates are denoted by λ and μ . The state of the system is specified by the pair $(L(t), S(t))$, i.e., one discrete and one continuous variable. The setting is very similar to that of a fluid queue driven by a Markov modulated rate process.

The evolution of the joint distribution of the state variables is governed by a system of coupled ordinary differential equations (ODEs) for the partial cumulative distribution functions $F_i(t, x) = P\{L(t) = i, S(t) \leq x\}$,

$$\begin{aligned} \frac{\partial}{\partial t} F_i(t, x) - \alpha(x-i) \frac{\partial}{\partial x} F_i(t, x) &= (\lambda F_{i-1}(t, x) - \mu F_i(t, x)) 1_{i>0} \\ &+ (\mu F_{i+1}(t, x) - \lambda F_i(t, x)) 1_{i<K}, \quad i = 0, \dots, K. \end{aligned} \quad (3.8)$$

An analytical stationary solution to these equations is found in a few special cases. A general stationary solution is not known and is believed not to have a simple form. Furthermore, it turned out that a direct numerical solution of the system of ODEs is unstable, and no simple remedy was found¹. Therefore, different alternative approximate ways for obtaining the stationary distribution are developed in the report.

Two of the methods consider the temporal behaviour of the state distribution. By Kolmogorov's theorem, starting from any initial distribution the system will eventually approach an equilibrium, i.e., integrating the equations in time is inherently stable. The first of the methods considers the evolution of the system in continuous time, whereas in the second approach an embedded system in discrete time is studied. A disadvantage of these methods is that the equilibrium is only approached asymptotically, and with a long averaging time the convergence is slow. The third method focuses directly on the equilibrium distribution but using an approximation. The method applies the stochastic discretization approach, where the deterministic evolution of the continuous variable is replaced by small stochastic transitions, thus allowing the use of standard methods of Markovian systems.

¹It has been later demonstrated by J.-M. Perkkiö (private communication) that a stable, fast-converging iterative numerical solution scheme does exist.

M/D/1/K vacation queue

In [7] a queueing model is adopted where voice packets are fed into a finite buffer and served by a server representing the output link. The aggregate voice traffic is modelled by a Poisson process. Due to the presence of best effort traffic and to the DiffServ-compliant non-preemptive priority scheduling, the operation of the server is considered in an exhaustive service and multiple vacation scenario. That is, the server serves voice packets until the buffer becomes empty. At the finishing instant of the service, if the server finds the queue empty, it takes a vacation. If there are still no voice packets in the queue when the server returns from its vacation, it takes another vacation, and so on. The vacations of the server correspond to the situation where the output link is occupied by the best effort traffic. The assumption of multiple vacations implies that the offered load of the best effort traffic is sufficiently high to utilize immediately the link capacity whenever no voice packet is present. The vacation time is assumed to be the time needed for transmission of a best effort packet with MTU size. In effect, a finite $M/D/1/K$ queue with exhaustive service and multiple server vacations is obtained.

The steady-state solution of this queueing model is obtained and useful quantities concerning packet loss probability and arbitrary percentile of delay are derived. The latter quantity is particularly valuable, because it stands for a statistical upperbound on the jitter, which is the main factor leveraging the perceived quality of voice connections.

BMAP/G/1 queue with feedback

Message transmission in wireless communication networks includes procedures of error correction at several layers of the protocol stack, e.g., at the data link (DLL) and transport layers. These procedures perform the repeated transmission of those protocol data units (PDUs) that were transmitted with errors. An adequate performance model of such procedures is determined by specific feedback queues.

In [147] matrix-geometric modeling techniques are used to analyze and calculate the performance characteristics of such a telecommunication channel where the probability of a corrupted transmission is fluctuating. Such situations typically occur during information transmission in mobile networks when the users cross cell boundaries and the interference conditions change

drastically.

First, the transmission process is modeled in terms of the BMAP/G/1 queue with feedback where the behavior of the input and the error probability depend on the state of a Markovian synchronous random environment. The latter describes the random changes of the interference conditions determining the success of a service completion. Here the probability of a repeated service, which can be interpreted as the error probability of a transmitted PDU, can change according to the state of that random environment. A general Batch Markovian Arrival Process (BMAP) describes the arrival stream of customers, modeling the batch arrival of radio blocks of a segmented message at the DLL. Applying the machinery of matrix-geometric methods, the resulting model is characterized by a discrete-time Markov chain with quasi block-Toeplitz structure embedded upon service completions. Then necessary and sufficient conditions for the existence of the corresponding steady-state distribution of the queue length at these embedded epochs are determined. Furthermore, the latter is characterized as the unique solution of a generalized variant of the Pollaczek-Khintchine equation using a generating-function approach. Finally, the stationary queue-length distribution at arbitrary epochs is determined and an algorithm for its calculation is sketched.

The model of [147] includes as special cases both Takacs' single-server feedback and a BMAP/G/1 queue operating in a synchronous random environment without feedback.

3.7 Queueing networks

The previous sections of this chapter deal with isolated queueing systems. In order to assess the performance of a communication network, however, it is also necessary to study networks of queues. This section overviews the COST-279 work on queueing networks. First, an approximate method to calculate end-to-end delay characteristics is presented. Next, a technique to determine the evolution of the characteristics of a traffic stream when it proceeds through a network is discussed. The latter can be useful to derive more accurate end-to-end performance characteristics.

End-to-end delay characteristics

The end-to-end delay is an important QoS parameter for real-time services. In [148], an analytical model to calculate end-to-end delays in packet networks is considered. The aim is to calculate the distribution of the end-to-end delay for a particular path consisting of a series of nodes. It is assumed that all the waiting times in the nodes in the end-to-end path are *statistically independent*; this is a key assumption to obtain the end-to-end delay by convolution. For queueing networks with FCFS queueing discipline this property only yields for the acyclic form of Jackson Networks (where a packet visits a node at most once). In [149], however, it is argued that if the load from a particular flow only is a small fraction of the total amount of traffic at a node and the input processes to the network are “smoother than Poisson” (i.e. with less variability) then the independent assumption will be quite reasonable and will represent a worst case scenario. Therefore, the M/G/1 queue is taken as the model to find the waiting time distribution in each node, and then the convolution is applied to obtain the end-to-end waiting time distribution.

If all nodes have identically distributed service times, the corresponding convolution may be substantially simplified and closed-form expressions are derived in terms of derivatives with respect to the load parameter. Special emphasis is put in [148] on the case with constant service times, since this is an important case for applications. Numerical results show that end-to-end delays in chains for up to 20 nodes may be analysed without numerical difficulties. It is also possible to extend some of the results to cover convolutions between equally loaded groups of queues with different service time distributions in each group.

Evolution of traffic characteristics

In [95], the evolution of the interarrival and interdeparture times between voice packets when they are proceeding through a number of network nodes has been investigated. Instead of separately specifying the characteristics of each individual source, the arrival process in a node is modelled as the superposition of a single tagged stream and an independent background process that aggregates the remaining traffic sources. Because the load of a single voice stream is assumed to be very low compared to the load of the aggregate traffic, the tagged voice packets can be represented as markers (packets with size zero). At the

entrance of each network node, one thus has the tagged traffic stream, namely the markers, and the background stream. The tagged marker stream is characterized by the consecutive interarrival times between the markers, which are assumed to be identically distributed. The background arrival process is described on a slot-per-slot basis according to a general i.i.d. process.

First, an expression for the probability generating function of the interdeparture times of the voice packets after one stage is established. The pgf of this interdeparture time is then used as the pgf of the interarrival times of the voice packets in the next stage, in order to assess the evolution of the interarrival-time characteristics throughout the network. Also, the pgf of the interdeparture times between three successive voice packets (in case two successive interarrival times may be dependent of each other) is calculated.

Chapter 4

Wireless Networks

Wojciech Burakowski, Andrzej Beben
Warsaw University of Technology, Poland

Contributors:

Samuli Aalto (Helsinki University of Technology, Finland), Hans van den Berg (TNO Telecom, Netherlands), Richard Boucherie (University of Twente, Netherlands), Llorenç Cerdà (Technical University of Catalonia, Spain), Sándor Imre (Budapest University of Technology and Economics, Hungary), Jorma Kilpi (VTT Information Technology, Finland), Michela Meo (Politecnico di Torino, Italy), Mihael Mohorčič (Jozef Stefan Institute, Slovenia), Dirk Staehle (University of Würzburg, Germany)

4.1 Introduction

Wireless mobile access has been developed into a most attractive way for users to connect to the Internet. However, users expect from the service providers new services with interesting content like news, local info, GPS info, and TV movies. To fulfil these expectations new requirements are placed on the network (including wireless access), particularly higher capacity of access links and end-to-end QoS guarantees at lower cost.

The statistics say that the number of users using mobile terminals will be growing each year and will in the near future overrun the number of fixed users. Therefore, intensive research and development efforts in the area of wireless and mobile networks are taking place, resulting in emerging new generations of networks. According to the current evolution scenario, in the near future the Third Generation (3G) systems, such as the UMTS and Code Division Multiple Access 2000 (CDMA2000), will substitute the existing Second Generation (2G) systems, namely GSM. The 3G systems are aimed at offering at least 144 kbps for outdoor (mobile) and 2 Mbps for indoor (fixed) links that are

considerable higher comparing with those offered by 2G systems, for instance only 9.6 kbps in GSM. On the other hand, WLAN technology, at the beginning treated as an extension of LANs, has recently experienced a phenomenal growth. Currently, the most popular 802.11 WLAN standard offers wireless access link rates up to 54 Mbps, with usage ranging from the office environment to public hot spots. Comparing WLAN and 3G systems, the WLANs have limited coverage but much higher link bit rates, whereas 3G systems have wider range and provide better mobility support. There is thus market pressure for the development of solutions allowing users to make seamless handovers between those two types of networks.

The discussed issue is how to set-up the most suitable connection depending on user application demands when an appropriately equipped terminal simultaneously has access to more than one wireless network, e.g. to WLAN and UMTS. For solving this problem, the concept of Always Best Connected (ABC) scenario is investigated.

Notice that the wireless links used in access are still of relatively lower bit-rates when compared to the wired links that are typical for the core. Therefore, for offering new services to the users, who usually require some end-to-end QoS guarantees, a new network architecture needs to be developed. This architecture should support QoS in both the core and the access networks, and user mobility. While there is a possibility to overprovision the core, for the wireless access networks some traffic control mechanism should be developed.

The research activities inside the COST 279 project on wireless and mobile networks mainly touch the selected issues of QoS traffic control, traffic measurements, modeling, performance evaluation, and resource management for GSM, UMTS, WLAN, and satellite technologies. These technologies present essential differences from each other, and as a consequence many proposed solutions are technology-specific. This chapter is therefore structured according to the main types of wireless networks.

4.2 GSM Networks

GSM is circuit-switched technology originally designed for handling voice calls. GPRS is an extension of GSM offering packet switching capabilities over the GSM channels. The technical specifications of GPRS were made in the 3rd Generation Partnership Project (3GPP) [150], and GPRS is currently

offered by many European GSM operators. GPRS uses the spare capacity left out by GSM voice or data calls. By default, GSM calls have priority in the use of radio channels. If available, the channels are reserved only temporarily when the terminal receives or transmits data. By using several time slots from the GSM frame, GPRS allows higher data rates compared to GSM data calls. These data rates and a quite short connection set-up time make the use of Wireless Application Protocol (WAP) services and Multimedia Messaging Service (MMS) cheaper and faster than by using GSM data calls. Moreover, connections to Internet are also possible.

When considering the performance issues of GPRS, the GSM radio channel is one of the main limiting factors, possessing an intrinsic structural delay and bandwidth variations and packet losses that are hard to deal with. The technical specifications give some indications for QoS, but these features are typically not taken in use.

4.2.1 Traffic Measurements and Modeling

In [72] and [84] are presented results obtained from `tcpdump` traces recorded between a GPRS backbone network and the Internet but before a NAT is made. The reported measurements are from about half a year after the operator had launched its commercial GPRS service. A GPRS session is defined as all packets (and their time stamps) with the same temporary IP address that the user is given at the time of the attachment to the GPRS network.

The main observations are that, during the measurement time, most GPRS sessions started during the working days and in working hours were in general roughly similar to low access speed dial-up sessions, and that GPRS session durations and session volumes seem to have heavy-tailed distributions. Also, the majority of the data transfer occurs typically in the beginning of the session, and the GPRS users typically detach from the GPRS network when they have finished active usage.

In paper [74] are studied distributional properties of GPRS/GSM session volumes and durations, by the application of a known statistical method of ascertaining about the underlying distribution to the data-analysis of GPRS data introduced in [72] and [84]. The statistical method is based on the concept of maximum correlation. This method is used because there are no *a priori* reasons to expect or prefer any particular distribution. The main result is that

these data sets exhibit a heavy-tailed nature. The Weibull distribution with the shape parameter between $0 < \alpha < 2$ is one of the possible models.

4.2.2 Resource Management

Resource management policies in GSM/GPRS cellular networks are studied in [151]. These systems basically offer two services: mobile telephony and wireless access to the Internet. The resource allocation policy should be carefully chosen so that telephony, from which operators are still getting most of their revenues, is not penalized and, at the same time, data service is provided with a good quality.

Three different alternatives are investigated. According to the *voice priority* channel allocation strategy, priority is given to voice in the access to radio channels. The second channel allocation policy is called *R-reservation*: it statically reserves a fixed number of channels to data services. Finally, the *dynamic reservation* strategy allocates channels to data whenever necessary, using the information about the queue length of GPRS data units within the base station. A threshold on the queue length is used in order to decide when channels must be allocated to data.

The schemes are analyzed by means of Markov chain models. From the results it is shown that voice priority cannot provide acceptable performance to data service, since whenever all the available channels are busy with voice connections, the data service undergoes service interruption. The R-reservation channel allocation policy overcomes this problem and drastically improves the performance of data service. The drawback of this scheme is that it subtracts resources from voice users, even when these are not needed for data, thus inducing an unnecessary performance degradation for voice services. The dynamic reservation scheme provides effective performance tradeoffs for data and voice services, with the additional advantage of being easily managed through the setting of the threshold value.

In [152] are considered adaptive dynamic channel borrowing strategies for wireless networks covering a road. In a Fixed Time Division Multiple Access (FTDMA) based network model, road traffic prediction models are used to characterize the movement of hot spots, such as traffic jams, and subsequently to predict volume of traffic load offered to the network. A dynamic upper bound on the capacity required to achieve a specified QoS level in the cells is computed. Restricting borrowing to neighbouring cells, to avoid ex-

cessive re-allocation of capacity, optimal channel borrowing strategies based on traffic movement and traffic density are given. These strategies can be characterized by a straightforward rule of thumb, of easy implementation: borrow capacity from the cell on the steeper side of the traffic peak. The simulation results under realistic traffic load conditions indicate a significant reduction of call blocking probabilities assuming the proposed optimal channel borrowing strategy.

The load balancing in cellular networks problem is discussed in [153]. The performance of both static and dynamic call routing policies for a simple model of two base stations with overlapping cells is compared. The new call originating from the overlapping area can be routed to either one of the two base stations. In a static policy the routing decisions are based only on the system parameters, i.e., arrival rates and mean holding times, assumed to be known, whereas a dynamic policy uses additional information about the system state, defined by the number of running calls in each station. Since the capacity of the base stations is finite, it may happen that a new call will be blocked. A method for constructing a reasonable routing policy (close to the optimal) is proposed.

The approach is based on the theory of Markov Decision Processes (MDP). The objective is to minimize the call blocking probability. In particular, the blocking performance of the policy obtained by the First Step of the Policy Iteration (FPI) algorithm is investigated. Under a randomized policy the two base stations can be modelled independently as Erlang loss systems, for which it is easy to determine the so called Howard's relative costs in each state. This makes the first step of the policy iteration algorithm of low complexity, and therefore applicable for large instances. In addition, it is easy to determine the optimal policy among the randomized policies (which is a subset of the set of static policies).

The idea is to approximate the optimal dynamic policy with the FPI policy. It is shown that any FPI policy fulfils two reasonable requirements: it is greedy, and of the threshold type. In addition, numerical experiments for small instances provide a deeper understanding of the FPI policy. Interestingly, it turns out that starting with the optimal randomized policy as the basic policy does not necessarily lead to the best performance of the FPI policy. Therefore, a heuristic rule is suggested for the basic policy. With such a choice, the FPI policy seems to be close to the optimal dynamic policy, and performs better

than the other considered policies.

4.3 UMTS Networks

UMTS is the 3G wireless network standard for providing different QoS services and operating with bit rates up to 2 Mbps. This is achieved by operation with Wideband Code Division Multiple Access (WCDMA) over the air interface.

Given the fundamental differences from GSM, the introduction of UMTS requires new paradigms in wireless network design. Recall that in GSM the capacity of a base station is determined only by the number of available frequencies and hence is independent of current network load. Given the number of frequencies available in a cell, the allowed network load follows directly from the Erlang-B formula, since the GSM network provides mainly voice calls. In contrast, the capacity of a Base Station (BS) in a WCDMA network is interference limited. On the uplink direction, the Multiple Access Interference (MAI) at a BS is caused by all the Mobile Stations (MSs) both from the given BS and from neighbouring BSs. On the downlink direction, the capacity is limited by the transmit power of the BS or by the interference it causes. The power control mechanisms in both link directions control the transmitted powers in such a way that, for each service, signals are received with nearly equal strength. A detailed examination of the interference on the uplink of a given cell is no straightforward task. Due to the universal frequency reuse in UMTS, all users, both in the considered cell and in the neighbouring cells, contribute to the total interference, thus influencing the link quality in terms of received bit-energy-to-noise ratio (E_b/N_0).

The planning of WCDMA networks consists of two aspects: the coverage planning and the capacity planning. In contrast to GSM, the coverage and the capacity cannot be considered as independent terms. In WCDMA a trade-off between the coverage area and the capacity of a BS exists. The more users are active at a BS, the larger is the MAI at the BS, and the higher are the transmit powers required by the MSs to fulfil their E_b/N_0 requirements. Additionally, due to the restriction of the MSs transmit power, the coverage area shrinks with an increasing number of users. Attaining a certain coverage area for a BS demands a limitation of the MAI, which can be done by admission control. The MAI level used as threshold for the acceptance of new calls determines

not only the coverage area, but also the capacity of the BS.

Another difference between GSM and WCDMA is the handover procedure. While GSM supports only hard handovers, where the connection to the new BS is established after terminating the one to the old BS (“break before make”), *soft handover* is performed in WCDMA. Here, the mobile assists in the handover process by measuring the pilot signals from the neighbouring BSs and storing those BSs with the strongest received signals in the Active Set (AS). The mobile then communicates with all BSs in the AS simultaneously (“make before break”). As a consequence, the MS receives multiple power control commands and adapts its transmission power on the uplink to the BS with the least requirement.

The introduction of 3G mobile communication systems also allows the service providers to offer a large variety of services, categorized in UMTS under the classes conversational, streaming, interactive, and background. While the conversational and streaming classes have a guaranteed bandwidth and delay, the interactive and background (best effort) classes consume the remaining system capacity. On the downlink this system capacity is limited by the BS transmit power, and on the uplink by the interference. As traffic in UMTS networks is expected to be asymmetric, with the bulk of it towards the MS, the downlink becomes the limiting link for best-effort traffic. In UMTS, best-effort traffic may either be carried on dedicated channels that are subject to rate control or on a shared channel. UMTS release 5 further standardizes the High-Speed Downlink Shared CHannel (HS-DSCH), which achieves data rates up to 10Mbps.

4.3.1 Admission Control and Capacity of WCDMA Systems

The key feature of WCDMA systems is that all users transmit in the same frequency band, their signals being separated by the use of orthogonal or pseudo-orthogonal codes. Except for the ideal case, when real orthogonal codes are used and no multi-path propagation occurs, a user sees the other users’ signal as interference. The total interference comprises the own-cell interference \hat{I}_{own} , the other-cell interference \hat{I}_{other} , and also the thermal noise \hat{N}_0 . The interference grows with the number of calls in progress and limits the capacity.

WCDMA Admission Control is performed on the basis of the measured

noise rise, defined as the ratio of the total interference \hat{I}_0 to the interference \hat{N}_0 of an unloaded (empty) system. The AC estimates the increase of the noise rise that would be caused by accepting a new connection and blocks it if the result exceeds a predefined threshold. While the noise rise is a value that is measured by a BS, it is not well suited for understanding the actual system load. A transformation of the noise rise yields the definition of the cell load η :

$$\text{Noise rise} = \frac{\hat{I}_0}{\hat{N}_0} = \frac{1}{\frac{\hat{N}_0}{\hat{I}_{own} + \hat{I}_{other} + \hat{N}_0}} = \frac{1}{1 - \frac{\hat{I}_{own} + \hat{I}_{other}}{\hat{I}_{own} + \hat{I}_{other} + \hat{N}_0}} = \frac{1}{1 - \eta} \quad (4.1)$$

A cell load equal to 1 defines the pole capacity of a WCDMA cell. On the arrival of new call submitted to service t the AC algorithm estimates the additional load α_t brought in by the call. This load is based on the negotiated traffic contract parameters, i.e., bit rate and maximum error rates. WCDMA AC consequently accepts an incoming connection if the estimated cell load η_{est} stays below the predefined threshold value η_{max} . The acceptance of a new call depends on both the own-cell interference and the other-cell interference. This feature explains why we speak of soft-blocking in WCDMA networks.

A time-efficient algorithm to compute blocking probabilities in a WCDMA network operating with several services is proposed in [154] under assumption that users submit calls to each service class accordingly to a Poisson process with exponentially distributed holding times. The load produced by a user is declared by the submitted QoS requirements, i.e., bit rate and target E_b/N_0 , and is called load per service. The user activity at an arrival instant is modeled by a Bernoulli random variable. The AC condition which originally relates to own- and other-cell interferences is transformed into own-cell load η_{own} and other-cell load η_{other} . The other-cell load is modeled as a lognormal random variable that is assumed to be independent of the own-cell load. This allows the derivation of the probability $\beta_t(\eta_{own})$ that a call of service t is blocked in a system state with own-cell load η_{own} . This probability is called the soft blocking probability and can occur in virtually every state depending on the other-cell load.

Figure 4.1 shows example blocking probability for a system operating with three services. The dotted lines show the approximated blocking probabilities and the solid lines correspond to simulation. One can see that the approximation yields accurate results in the presented range of scenarios, possessing




Figure 4.1: Uplink Blocking probabilities for three services with 12.2kbps, 64kbps, and 144kbps

strongly varying other-cell interference levels and different user activities of 0.45, 0.3, and 0.8 for the three services. These are the especially interesting cases, since with deterministic other-cell interference and Always-ON users the proposed analysis yields exact results.

An alternative way to perform AC in WCDMA networks is discussed in [12]. The AC algorithm combines the dynamic optimization of the Chernoff bound instead of the well-known static effective bandwidth concept. A very important advantage of the proposed method is its dynamic behaviour that, in contrast to the traditional static effective bandwidth methods, allows resilient adaptation to the continuously changing network parameters. The proposed algorithm is able to adapt dynamically to an ever-changing radio environment, and provides a trade-off between decision efficiency and complexity. The proposed AC method is investigated under ON/OFF traffic sources and lognormal fading channels.

In [155] is analyzed downlink feasibility of call assignments to base stations in a heterogeneously loaded, linear UMTS network. This model is motivated by the situation along a highway where, due to traffic jams (“hot spots”), the load of the cells is not distributed evenly along the road. By dividing the

area into small segments, the power requirements are characterized via a matrix representation that separates user and system characteristics. A closed-form expression of the Perron-Frobenius (PF) eigenvalue of that matrix is obtained, providing a quick assessment of the feasibility for each distribution of calls over segments (and for each assignment of calls to base stations). In particular, the PF eigenvalue is almost linear in the number of calls per segment, and thus provides a kind of “effective interference characterization of downlink feasibility”. The results allow for a fast evaluation of outage and blocking probabilities and may be used for call acceptance control.

In [156] is examined the evaluation of interference in direct sequence spread spectrum mobile communication systems. Interference conditions in both the uplink and downlink transmission are discussed. For calculation of interference conditions, a statistical method based on the average emission power of a mobile station in a cell is proposed and compared with one based on calculation of multiple cell interference reduction factors. The study shows that these two methods produce similar results.

4.3.2 Soft Handover

In CDMA systems, during the soft handover process MSs are connected, not only to one, but to several BSs. An MS moving in an area with several BSs has an Active Set that changes dynamically and is determined by the pilot signal transmitted by every BS. An MS detects the BS with the strongest received pilot signal and also those BSs with a signal strength not more than the *reporting range* below the strongest signal. All these BSs form the Active Set of an MS.

On the uplink, all BSs in the Active Set receive the frames transmitted by the MS (*site diversity*) and transfer them to the Radio Network Controller (RNC). There, all frames are checked for errors and only if all of them are erroneous a frame error occurs (*selection diversity*), see Figure 4.2. The RNC evaluates the resulting frame error rate and adapts the target E_b/N_0 in the outer loop power control. This target E_b/N_0 is signaled to all BSs in the Active Set, which try to adjust the transmission power of the MS to this value according to the inner loop power control. The MS receives power control signals from all BSs in the Active Set and increases its power only if all BSs signal *power up*. Otherwise, if one or more BSs signal *power down*, the MS obeys the latter command. On the uplink, soft handover leads to a reduction of the required transmission powers and consequently to less interference in




Figure 4.2: Diversity effects due to soft handover in WCDMA networks

the system, from which an increased system capacity results. The benefits of soft handover on the uplink include automatic load balancing, a target E_b/N_0 reduction, and increased robustness against fading. The automatic load balancing and the robustness against fading result from the inner loop power control, which “selects” the best BS on slot level, i.e., 15 times per 10ms interval. The best BS is not solely determined by the propagation loss, but also depends on the interference levels at the BSs in the Active Set. MSs at the border of a highly loaded BS may be power-controlled by a nearby BS with less load, even though the propagation loss to this BS is higher. Thus, the load is shifted from BSs with high load to BSs with lower load. This effect is investigated in [157] using Monte Carlo simulation techniques. A series of snapshots of a UMTS network with 39 BSs in a hexagonal grid is generated for homogeneous and non-homogeneous spatial traffic distributions. The soft handover gain, defined as the reduction of the mean interference due to soft handover, grows in the homogeneous scenario with the offered load, and reaches a maximum of about 3dB. In non-homogeneous scenarios even higher soft handover gains are obtained, since there is more potential for load balancing. The outer loop power control further increases the soft handover gain by decreasing the target E_b/N_0 value such that lower received powers are required at the BSs. This gain is additional to the benefit obtained by selecting the best BS, as still all BSs in the Active Set try to adjust the MSs transmit powers to the same target value. The effect of the outer loop power control is investigated in [158]. In this paper soft handover gains up to 8dB are reported in homogeneous net-

works with high traffic density.

4.3.3 Impact of mobility on UMTS network performance

The mobility of terminals in a UMTS network strongly influences network capacity. On the one hand, this impact stems from the more severe Signal-to-Interference-Ratio (SIR) requirements that apply in case of higher velocities due to the combined effects of multipath propagation, Doppler shifts, and power control imperfections. On the other hand, an increased mobility requires a higher level of radio resource reservation regarding handovers, in order to keep the call dropping probability below a prespecified target value. As a consequence, fresh call blocking increases, inducing a need for denser site planning. In [159] an analytical approach is presented to evaluate the impact of these two mobility-related aspects on network planning, performance, and investment costs. The principal strength of the approach lies in the model being simple enough to allow a computationally relatively inexpensive performance evaluation and optimization, yet being sufficiently realistic to provide valuable qualitative insight for network planning purposes.

The fresh call blocking probabilities and call dropping probabilities are evaluated using a two-dimensional time-continuous Markov chain with a state described by the number of users in the considered cell and in the surrounding cells. The interference of a user in a surrounding cell is modelled by a Gaussian random variable including terminal location and call activity. By demanding a given outage probability, the soft capacity of a WCDMA cell is transformed into a fixed set of accessible states. The velocity of the users is incorporated into the different SIR requirements and cell residence times. Furthermore, the call admission control for fresh calls reserves a certain radio resource for handover calls. The amount of reserved radio resource and the cell radius are optimized in order to minimize network investment costs in terms of BSs per covered area, while meeting the target call blocking and dropping probabilities. The primary conclusions from the numerical examples are that (i) the impact of terminal velocity on the optimal cell radius, and therefore on the investment costs, can be quite significant (potentially up to a factor of 2), (ii) the deployment of a radio resource reservation scheme can indeed be effectively utilized to reduce call dropping and investments costs, and (iii) planning

a UMTS network using inaccurate estimates of terminal velocity can lead to unacceptable blocking and, in particular, dropping probabilities.

In [159] the influence of mobility on the radio resource reservation for handover calls is demonstrated using a simple mobility model, with cell residence times with exponential distribution and expectation dependent on the mobile's velocity. In [160] a new mobility model is introduced in order to predict the number of terminals in each cell and to improve the effectiveness of CAC algorithms. This model is studied in the context of other mobility models reported in the literature, namely random walks and Markov models.

Two-dimensional Markov models are very complex, because of the need to keep state for the six adjacent cell directions. The goal to define a simple yet appropriate model is achieved by extending the one-dimensional model to a two-dimensional model that retains its simplicity by limiting the possible states of the mobile user. The main idea of the model is to separate the neighbouring cells into two groups according to the typical user movement direction. The authors have calculated the number of mobile terminals for future time periods relying on the random walk model and their modified Markov mobility model. The calculation of the future number of mobile terminals in cell k uses the concept of a ring, defined as the cells surrounding cell k . This concept simplifies the calculations, because the interest rests only on the number of users arriving to or leaving a given ring during a time period, whereas internal movements (inside the rings) remain unconcerned. The goal is to predict the number of users in cell k in time slot t_{i+1} and t_{i+2} , based on the number of users in cell k , the first ring, and the second ring at time slot t_i . The authors have compared the accuracy of different mobility models by simulation. The results and equations obtained can be utilized in resource-reservation based AC algorithms, paging algorithms, etc. The parameter of the random walk model is easily determined from measurements of the actual number of handovers in the network. The parameters of the modified Markov model can also be calculated from handover measurements. In addition, the directions of the handover events and the time interval between the handover events are needed. Predictions based on the extended Markov model prove to be as precise as the handover vector random walk method, but the required calculations and equations are much simpler. The results show that the accuracy of the prediction depends on the range of the forecast. The two-ring concept takes into consideration the users moving at higher speed, therefore provides more accurate

information for the AC algorithm than the ring based forecast. It is left open for future work the determination on the optional forecast distance as function of the mobile terminal's speed in the network.

4.3.4 Scheduling of Packet Data Traffic in UMTS Networks

Wireless data transfer is undisputedly a major driver for the deployment and anticipated success of third generation mobile networks. As the foreseen (data) services greatly differ in their traffic characteristics and QoS requirements, a number of distinct transport channels have been specified to accommodate these services efficiently. In the downlink, commonly the focus of the attention in light of the generally expected strong up/downlink data traffic asymmetry, the Dedicated CHannel (DCH), the Forward Access CHannel (FACH), and the Downlink Shared CHannel (DSCH) are standardized. Specifically designed for delay-sensitive services or services with stringent throughput requirements, the DCH is a bit pipe assigned exclusively to a single mobile station and has the advantage of fast closed-loop power control and macro-diversity. On the other hand, the FACH and DSCH may be shared by multiple mobiles. The FACH is typically used for the transfer of relatively small data chunks, without the advantages of closed-loop power control and macro-diversity. Medium to large data transfers, particularly of bursty character (e.g., TCP/IP flows), are most efficiently conveyed on the DSCH, as it enjoys the advantages of closed-loop power control by maintaining a low bit rate Associated DCH (A-DCH) for signalling purposes. A principal advantage is the enhanced efficiency of channelization code usage. Since data is multiplexed on the DSCHs, the use of soft handover (or macro-diversity) is rather complicated from an implementation viewpoint and therefore not standardized. UMTS release 5 further standardizes the HS-DSCH, which achieves data rates up to 10Mbps by a variety of enhanced technologies including (i) higher order modulation together with fast link adaptation, i.e., adaptive coding and modulation, optimized for channel conditions, (ii) fast scheduling centered at the Node-B rather than at the RNC, with a proposed smaller transmission time interval, in order to reduce delays and facilitate better tracking of the channel variations (e.g., equal to the 0.67ms slot duration), and (iii) fast cell selection in order for an MS to continuously select the serving cell with the best radio conditions, important

in light of the unavailability of macro-diversity on the HS-DSCH.

In [161] the authors present a semi-analytical performance evaluation of the DSCH, investigating the influence of A-DCHs and interference from neighbouring cells on the downlink transfer rates. The analysis of a UMTS network with N cells consists of two stages. In the first stage, the cell-wise outage probabilities for a given number of calls per cell are obtained using Monte Carlo simulation. The second stage captures the traffic dynamics of call arrivals and terminations in an N -dimensional irreducible continuous-time Markov chain. A system state is specified by the number of calls per cell, and the outage probabilities obtained in the first stage are integrated by reduced arrival rates. Furthermore, for each state and cell an effective throughput per call is determined as function of the frame error rate.

The influence of A-DCHs and other-cell interference is demonstrated by simple experiments. The data transfer rates of networks with one, two, and three Node-Bs are evaluated with and without A-DCHs. The experiments show that a heavier data traffic load implies both a greater competition for DSCH resources, and thus longer transfer delays, and a higher interference level due to the greater number of A-DCHs that must be maintained for signalling purposes. This latter effect causes a higher frame error rate and thus a lower effective aggregate DSCH throughput. Hence, the greater the demand for service, the smaller the aggregate service capacity.

In [162] the authors investigate the performance of an integrated services UMTS network with speech calls on DCHs and data calls on an HS-DSCH. Different scheduling schemes for micro- and macro-scheduling are compared. *Macro-scheduling* adjusts the power of the HS-DSCH to the varying traffic conditions at a time scale of seconds and has inter-cellular scope. *Micro-scheduling* time-multiplexes the data flows within each cell in order to optimize resource efficiency according to varying channel conditions, while satisfying the call's QoS requirements and providing some sort of fairness. Two different scheduling disciplines are considered for micro-scheduling: *power-fairness* and *rate-fairness*. Power-fairness means that within one cell the same power is spent for all users, and thus users with worse channel conditions experience a lower rate. Rate-fairness means that all users within one cell obtain the same rate, but more power is spent for users with worse channel conditions. Macro-Scheduling can be either *adaptive* or *fixed*. Adaptive means that the power allocated to the HS-DSCHs of the different cells optimizes the data

throughput while maintaining the C/I requirements of the speech users.

The different disciplines for micro- and macro-scheduling are compared using Monte Carlo simulation techniques. For a series of snapshots, the outage probability for speech calls and the achieved rate for data users are evaluated as performance criteria. The principal objective of the experiments is to demonstrate the potential performance enhancements of adaptive macro-scheduling. The results show that indeed adaptive scheduling leads to an improved QoS for both speech and data services. Speech calls enjoy the most significant performance gain due to the strict prioritization strategy. For the fixed macro-scheduling scheme there is a trade-off between the QoS for speech and data users. More power for the HS-DSCH leads to a higher transfer rate for data users, but also to a higher outage probability for speech users. Power-fair micro-scheduling leads to a higher expected throughput per data call than rate-fair micro-scheduling. However, the expected throughput for users at the cell border is significantly smaller than for users near the center of a cell.

4.4 Wireless LANs

Recently, WLANs have gained a prominent role in the telecommunications environment because of their support of seamless access of portable devices, like laptops and electronic organizers, to the Internet. As the Internet becomes a multi-service network, one can also expect from WLANs the internal capabilities for providing guarantees on some QoS parameters (delays, packet loss). WLAN standards, like HIPERLAN and IEEE 802.11, are mainly concentrating on the physical, data link, and MAC layers. However, for supporting user mobility in a wide range environment (outdoors), enhancements at the network layer, e.g., mobile IP, are required. Physical layer issues need to be taken into account because interference caused by other systems (WLANs, Bluetooth) or by noise strongly influences the available capacity of the system. From a QoS point of view, efficient scalable MAC protocols that can handle traffic with different QoS requirements are expected. The IEEE 802.11e standard [163] attempts to fulfil these demands, but only with moderate success. In addition, roaming and handover have to be considered, since a significant amount of delay and packet loss can be traced back to these events.

The activities of the COST 279 project in the area of WLANs have mostly focused on performance analysis, handover issues, and reliability aspects.

4.4.1 Performance analysis

An integrated packet/flow level model suitable for analyzing flow throughputs and transfer times in IEEE 802.11 WLANs is proposed in [99]. The packet level model captures the statistical characteristics of the transmission of individual packets at the MAC layer, while the flow level model takes into account the system dynamics due to the initiation and completion of data flow transfers. The model is a processor sharing type of queuing model, reflecting the IEEE 802.11 MAC design principle of distributing the transmission capacity fairly among the active flows. The resulting integrated packet/flow level model is analytically tractable and yields a simple approximation for the throughput and flow transfer time. Extensive simulations show that the approximation is very accurate for a wide range of parameter settings. In addition, the simulation study confirms the attractive property following from the approximation that the expected flow transfer delay is insensitive to the flow size distribution (apart from its mean).

4.4.2 Handover issues

In [164] is discussed a queuing network mathematical model for the analysis of the low latency handoff schemes related to those of [165, 166, 167, 168, 169]. Consider the network architecture depicted in Figure 4.3. For computational tractability all routers are modeled as simple M/M/1 queues. The authors focus on the Pre-Registration handoff scheme, the model corresponding to the Post-Registration scheme being presented in [168].

Consider a Mobile Node (MN) moving from the old Foreign Agent (oFA) to the new Foreign Agent (nFA) and suppose an overlapping area between the two subnetworks. Assume that the L2 handoff starts when the MN enters the overlapping area, and denote this time instant by t_0 . Let the variables D_{ST} , D_{LD} , and D_{LU} define the time needed, since t_0 , to generate the layer 2 triggers L2-ST (*Source Trigger*), L2-LD (*Link Down*), and L2-LU (*Link Up*), respectively. These are constant positive values such that $D_{ST} < D_{LD} < D_{LU}$. Furthermore, the response time of any router A is denoted by R_A . The time instant the Gateway Foreign Agent (GFA) starts forwarding to the nFA, instead of to the oFA, the packets destined to the MN is denoted by t_1 and given by $t_1 = D_{ST} + R_{oFA} + R_3 + R_{nFA} + R_2 + R_{GFA} + \text{fixed delays}$ (see Figure 4.3). This expression is a sum of exponentially distributed variables




Figure 4.3: Network architecture.




Figure 4.4: Delay distribution.

and constants. Packets will be routed via the oFA or via the nFA according to whether they arrive at the GFA before or after t_1 .

As an example, Figure 4.4 shows the expected number of CBR packets that are dropped due to expiration of the playout time or the absence of a buffer, as function of the playout time. Both the results for Pre-Registration and Post-Registration are shown, for two different values of the time between the LD and the LU triggers. The timing of the LD trigger is set to 60 ms. Note that these curves tend to the expected number of lost packets due to the absence of buffer capacity when the playout time tends to infinity. It can be seen that Pre-Registration implies more losses than Post-Registration, while the average delay for packets that are not lost is slightly larger for the Post-Registration scheme. The latter follows from the fact that packets using the bi-directional edge tunnels have a longer delay. More packets are lost when the time between the LD and the LU trigger increases, so more buffer capacity would be needed to avoid losses. In the paper, issues related with handoff implementations over IEEE 802.11 are also discussed.

At present, new issues of high-speed network research emerge from the convergence of wireless and wired next generation networks and the planned deployment of new integrated multi-media and Web services. In [80] the authors have sketched a service architecture where adaptive applications are running above a distributed service management and QoS-control layer. These layers operate, in turn, on top of the transport and network layers with their corresponding functionality of resource allocation and control, and of mobility

management. They have described a LINUX implementation of the proposed concept in an IEEE802.11b compatible wireless LAN with Mobile IP (MIP) as network layer using a unified programming model, and have investigated the efficiency of the data transport by the TCP and UDP protocols. By means of appropriate measurement tools is illustrated the impact of the changing transmission quality, the resulting error recovery of the wireless data link layer, and the roaming between different basic service areas on the dynamics of TCP flow control and on the resource-reservation process taking into account realistic emission patterns of the IP frames. Moreover, the most effective and sensitive control information of the adaptation process is identified.

As a summary, the study identifies some important issues regarding the interworking between resource reservation, QoS-control, mobility-, and security-management that require an improved, more efficient solution in the near future. From experiments it is concluded that an interworking between the TCP flow-control and the behavior of the wireless Data Link Control (DLC) layer should be organized to improve the throughput and delay characteristics of the data flows while moving and roaming, e.g., by exchanging management and control information that optimizes an ECN-based TCP flow control in a way similar to the Available Bit Rate (ABR) scheme in ATM. The use of improved TCP variants, e.g., those that freeze and recover the congestion-window state after a handoff, that use TCP control block statistics, that mimic a loss-less PDU transfer at the data link layer like snoop, or that distinguish different sources of packet loss may be another alternative.

Despite the technical shortcomings of the implementation, it is further demonstrated that micro-mobility is not supported in an adequate manner by the current functionality of MIP. Cellular IP, route and address caching techniques, improved movement detection by eager or hinted cell switching, and fast or two-phase handover schemes may partially resolve the observed difficulties. One of the most important issues is the lack of a co-ordinated interworking between MIP, the resource management, and the handover management. This problem has to be solved in an efficient way, otherwise the performance of the handoff process will deteriorate considerably and TCP flow-control may time out, causing massive performance degradation.

In conclusion, the realization of concept and the measurements clearly reveal the potential and drawbacks of the mobile-aware approach and provide new insight on the relationship of the signal-to-interference ratio, the delay-

loss characteristic of flows, and the TCP behavior, as well as the adaptation processes. The results may be used to develop improved handoff and TCP flow-control mechanisms as well as improved mobile-aware applications for an efficient transport of real-time multi-media and interactive Web services in next generation wireless networks.

The study in [170] refers to a WLAN network based on ATM technology. The problem of extending the connection-oriented ATM technique to mobile environments is that the users may change access point to the network during an on-going connection (handover). As cellular networks are using smaller cells to increase frequency reusability and system capacity and to decrease used powers, the rate of handovers increases and can cause the overload of the network management if all the handovers are followed by a connection set-up procedure. The above problem could be solved with more efficient exploitation of existing resources implemented with the use of suitable handover procedures. Solving the problem by reserving channels also in the neighboring cells leads to inefficient exploitation of resources, since the mobile user entering only one of the neighboring cells makes the reservation of free capacity in the other cells unnecessary. Nevertheless, the QoS has to be guaranteed for the connection in case of handover, which otherwise cannot be ensured if the mobile user enters an overloaded cell. In addition to these problems, mobile systems with non-voice type transmission also require considerable and variable bandwidth. There are several approaches to fight this problem, such as use of the virtual connection tree, shadow cluster, umbrella cell, location prediction, and handover supporting routing algorithm. The radio link related part of the handover is often in focus of research activities. Several publications address this problem, and many solutions have been published based on the idea of channel assignment and bandwidth reservation to support handover. The novelty of the scheme proposed in [170] lies in the fact that it focus not only on the handover problem at the radio interface, but also involves the wired network and extends the routing with a sophisticated new Location Prediction algorithm to speed up the handover process. A short survey is given about solutions for the support of efficient handover (virtual connection tree, shadow cluster, umbrella cell, location prediction or routing), and a comparative study is made that shows that they can fulfil the QoS requirements. A new handover supporting routing method is also presented in the paper, and its performance investigated by means of computer simulation.

4.4.3 Reliability aspects of mobile IP

The mobility provided by the Internet Engineering Task Force (IETF) Mobile IP cannot properly fulfil the requirements of an always-on scenario; hence micro mobility is needed to extend macro mobility. There are several micro mobility protocol recommendations introduced in the literature, most of them based on a physical or logical tree topology network. The most important and severe weakness of the tree topology is its poor reliability. In [171] the reliability of mobile IP is investigated. For this purposes a graph model of mobile IP networks is introduced where the edges of the graph represent the links connecting the nodes. In the model only links break down, the nodes being totally reliable. The links have two states, up (working), and down (broken). In this simple model all the links have an independent (and very low) probability of being in the down state. Reliability means that faults in the system do not degrade the performance of the system too much. To formalize this statement a performance function is introduced as a reliability measure, representing the performance of the system in a state. Maximum performance is the value of the performance function in the faultless state of the system. If the performance of the system in a given state is divided by the maximum performance we get the relative performance. The reliability measure is the following: The performance of the network in a given state is the number of base stations that can reach the backbone. For general tree topology micro mobility networks a recursive algorithm is proposed for the calculation of the exact distribution function of the relative performance. This algorithm works for any kind of tree topology micro mobility network. Using the same reliability measure, the algorithm can be extended for other network topologies.

4.5 Satellite communication

In [93] is described a traffic flow model developed for the use with an Inter Satellite Link (ISL) network simulator to modulate an arbitrary packet generator on the origin satellite (satellite serving the source terminal) and thus generate data packets that are then routed through the network to the destination satellite (satellite serving the destination terminal). The main objective of this work is the development of a complex simulation framework for testing and analyzing the performance of an arbitrary adaptive routing algorithm, hence it

was built in a modular approach.

It is shown that the assumed traffic flow model is suitable for performance evaluation of global satellite networks with an arbitrary satellite constellation. However, because of specific implementation issues, the study focused on a Low Earth Orbit (LEO) satellite system based on an inclined (delta pattern) constellation of 63 satellites in 7 orbital planes. In this configuration, the orbital planes have an inclination angle of 48° with respect to the equatorial plane, and satellites orbit the Earth at 1400 km with an orbit period of 114 minutes. The ISL network assumes permanent topology with two intraplane and two interplane ISLs per satellite.

In [93] is addressed the issue of per-hop packet routing, particularly well suited to the regular mesh topology with multiple alternative paths of connectionless ISL networks. In particular, it is proposed the use of Traffic Class Dependent (TCD) routing, capable to find suitable paths that can accommodate different types of services using different optimization criteria. However, similarly to conventional single-service routing procedures deployed in packet networks, TCD routing still does not guarantee the provision of any minimum requirements as in QoS routing. Three different representative traffic classes with diverse requirements have been introduced to evaluate the performance of the TCD routing procedure in the ISL network, each being routed according to its particular optimization criteria.

The performance of the proposed TCD routing procedure in different traffic load scenarios is evaluated using a simulation model of the ISL network and compared to the performance of a simple single-service routing, which makes no distinction between packets belonging to different traffic classes. Also, results obtained with empirical traffic are compared with those obtained using equivalent Poisson traffic. In both cases, homogeneous and non-homogeneous traffic flows between satellites are considered. Simulation results are presented as a trade-off of two performance measures, average relative packet delay deviation and average normalised data throughput.

List of Abbreviations

3GPP	3 rd Generation Partnership Project
ABC	Always Best Connected
ABR	Available Bit Rate
AC	Admission Control
ADSL	Asymmetric Digital Subscriber Line
AF	Assured Forwarding
AMP	Adaptive Multi-Path Routing
AMS	Anick-Mitra-Sondhi
AQM	Active Queue Management
AQUILA	Adaptive Resource Control for QoS using an IP-based Layered Architecture
ARQ	Automatic Repeat Request
AS	Active Set
ATM	Asynchronous Transfer Mode
b2b	border-to-border
BBB	Border-to-Border Budget
BD	Birth and Death Process
BMAP	Batch Markovian Arrival Process
BMBF	German Ministry for Education and Research

BS	Base Station
BU	Burstification Unit
CAC	Connection Admission Control
CBR	Constant Bit Rate
CCD	Complementary Cumulative Distribution
CDMA2000	Code Division Multiple Access 2000
CLT	Central Limit Theorem
DB	Delay Budget
DCH	Dedicated CHannel
DDR	Decompose-Design-Reassemble
DLC	Data Link Control
DLL	Data Link Layer
DSCH	Downlink Shared CHannel
EB	Egress Budget
ECMP	Equal Cost Multi-Path
ECN	Explicit Congestion Notification
EF	Expedited Forwarding
ELB	Egress Link Budget
EMW	Elwalid-Mitra-Wentworth
FACH	Forward Access CHannel
FBM	Fractional Brownian Motion
FCFS	First-Come-First-Served
FDL	Fiber Delay Line
FEC	Forward Error Correction
FFM	Fluid Flow Model
FIFO	First In First Out
FPI	First Step of the Policy Iteration
FTDMA	Fixed Time Division Multiple Access

FTP	File Transfer Protocol
GFA	Gateway Foreign Agent
GMPLS	Generalized Multi-Protocol Label Switching
GPRS	General Packet Radio Service
GPS	Generalized Processor Sharing
GSM	Global System for Mobile
HOL	Head-of-Line
HS-DSCH	High-Speed Downlink Shared CHannel
HTTP	HyperText Transfer Protocol
i.i.d	independent and identically distributed
IB	Ingress Budget
IETF	Internet Engineering Task Force
ILB	Ingress Link Budget
IP	Internetwork Protocol
IS-IS	Intermediate System to Intermediate System
ISL	Inter Satellite Link
ISP	Internet Service Provider
IST	Information Science and Technology
KING	Key components for the Internet of the Next Generation
LAC	Link Admission Control
LAN	Local Area Network
LB	Link Budget
LEO	Low Earth Orbit
LISDLG	Limit of the Integrated Superposition of Diffusion processes with Linear differential Generator
LRD	Long Range Dependence/Dependent
LSP	Label Switched Path
MA	Measurement Area

MAC	Medium Access Control
MAI	Multiple Access Interference
MBAC	Measurement Based Admission Control
MDP	Markov Decision Process
MHS	Message Handling Service
MIP	Mobile IP
MMFP	Markov Modulated Fluid Process
MMPP	Markov Modulated Poisson Process
MMRP	Markov Modulated Rate Process
MMS	Multimedia Messaging Service
MN	Mobile Node
MP	Measurement Point
MPEG	Moving Pictures Expert Group
MPLS	Multi-Protocol Label Switching
MRDV	Multipath Routing with Dynamic Variance
MS	Mobile Station
MTU	Maximum Transfer Unit
NAC	Network Admission Control
NAT	Network Address Translation
nFA	new Foreign Agent
NGN	Next Generation Network
NJ	Negligible Jitter
NRT	Non Real Time
OBS	Optical Burst Switching
OD	Origin-Destination
ODE	Ordinary Differential Equation
oFA	old Foreign Agent
OSPF	Open Shortest Path First

OUP	Ornstein-Uhlenbeck Process
P2P	Peer-to-Peer
PCAP	Packet CAPture
PCBR	Premium Constant Bit Rate
pdf	probability density function
PF	Perron-Frobenius
PFS	Priority Forcing Scheme
pgf	probability generating function
PLR	Packet Loss Rate
PMC	Premium Mission Critical
pmf	probability mass function
PMI	Passive Measurements Infrastructure
PMM	Premium Multi Media
POP	Point of Presence
PS	Processor Sharing
PVBR	Premium Variable Bit Rate
PVC	Permanent Virtual Circuit
QBD	Quasi-Birth-Death
QoS	Quality of Service
RAM	Random Access Memory
RED	Random Early Detection
REM	Rate Envelope Multiplexing
RKHS	Reproducing Kernel Hilbert Space
RMON	Remote Monitoring
RNC	Radio Network Controller
RSM	Rate Sharing Multiplexing
RSVP	ReSerVation Protocol
RT	Real Time

RTT	Round Trip Time
ScRM	Scalable Reliable Multicast
SIMA	Simple Integrated Media Access
SIR	Signal-to-Interference-Ratio
SiRM	Simplified Reference Model
SMA	Super Measurement Area
SMTP	Simple Mail Transfer Protocol
SNMP	Simple Network Management Protocol
SNR	Signal-to-Noise Ratio
TCD	Traffic Class Dependent
TCP	Transmission Control Protocol
TDAC	Traffic Descriptor based Admission Control
TM	Traffic Matrix
UDP	User Datagram Protocol
UMTS	Universal Mobile Telecommunications System
WAN	Wide Area Network
WAP	Wireless Application Protocol
WCDMA	Wideband Code Division Multiple Access
WLAN	Wireless Local Area Network
WVPing	Wave-Ping

Bibliography

- [1] Joachim Charzinski, Cornelis Hoogendoorn, Karl Schrodi, Christian Winkler, and Manfred N. Huber. Towards Carrier-Grade Next Generation Networks. Technical Report 279TD(03)040, COST-279, 2003. [cf. COST-279, [TD\(03\)040](#)]. [14](#), [26](#)
- [2] A. Bak, W. Burakowski, F. Ricciato, S. Salsano, and H. Tarasiuk. Traffic Handling in AQUILA QoS IP Networks, Quality of Future Internet Services. In M. I. Smirnov, J. Crowcroft, J. Roberts, and F. Boavida, editors, *Quality of Future Internet Services: Second COST 263 International Workshop, Qofis 2001, Coimbra, Portugal, September 24-26, 2001, Proceedings*, volume Volume 2156/2001 of *Lecture Notes in Computer Science*, pages 243–260. Springer-Verlag Heidelberg, January 2001. [cf. COST-279, [TD\(01\)017](#)]. [16](#), [24](#)
- [3] K. Lindberger. Dimensioning and design methods for integrated ATM networks. In *Proc. of 14-th International Teletraffic Congress, Antibes, 1994*, 1994. [17](#)
- [4] C. Brandauer and P. Dorfinger. An Implementation of a Service Class Providing Assured TCP Rates with the AQUILA Framework. In W. Burakowski, B. F. Koch, and A. Beben, editors, *Architectures for Quality of Service in the Internet 2003*, number 2698 in *Lecture Notes in Computer Science*, 2003. [17](#)
- [5] W. Burakowski and H. Tarasiuk. Admission Control for TCP Connections in QoS IP Network. In C. Chung, C. Kim, W. Kim, T. Ling, and K. Song, editors, *Web and Communication Technologies and*

- Internet-Related Social Issue - HSI 2003: Second International Conference on Human.Society@Internet, Seoul, Korea, June 2003, Proceedings*, volume 2713/2003 of *Lecture Notes in Computer Science*, pages 283–293. Springer-Verlag Heidelberg, January 2003. [cf. COST-279, [TD\(02\)029](#)]. [17](#), [20](#), [21](#)
- [6] M. Dabrowski, G. Eichler, M. Fudala, D. Katzensgruber, T. Kilkanen, N. Miettinen, H. Tarasiuk, and M. Titze. Evaluation of the AQUILA Architecture: Trial Results for Signalling Performance, Network Services and User Acceptance. In W. Burakowski, B. F. Koch, and A. Beben, editors, *Architectures for Quality of Service in the Internet 2003*, number 2698 in *Lecture Notes in Computer Science*, 2003. [17](#)
 - [7] H. T. Tran, T. Ziegler, and F. Ricciato. QoS Provisioning for VoIP Traffic by Deploying Admission Control. In *LNCS 2698, proceedings of Workshop on Architectures for Quality of Service in the Internet, Art-QoS*, pages 139–153, March 2003. [cf. COST-279, [TD\(02\)037](#)], [[172](#)]. [18](#), [95](#), [145](#)
 - [8] T. Bonald, A. Proutier, and J. W. Roberts. Statistical Performance Guarantees for Streaming Flows using Expedited Forwarding. In *Proceedings of IEEE INFOCOM*, volume 2, 2001. [18](#)
 - [9] Marek Dabrowski and Wojciech Burakowski. Assessment of Token Bucket Parameters by On-Line Traffic Measurements. In *10th Polish Teletraffic Symposium*, September 2003. [cf. COST-279, [TD\(03\)044](#)]. [19](#), [52](#)
 - [10] A. Elwalid, D. Mitra, and R. H. Wentworth. A new approach for calculating buffers and bandwidth to heterogenous, regulated traffic in an ATM node. *IEEE Journal on Selected Areas in Communications*, 13(6), August 1995. [19](#)
 - [11] Michael Menth and Oliver Rose. Performance Tradeoffs for Header Compression in MPLS Networks. In *10th International Telecommunication Network Planning Symposium*, pages 503–508, Munich, Germany, 2002. [cf. COST-279, [TD\(02\)031](#)]. [19](#), [20](#)
 - [12] S. Imre, P. Petras, and R. Tancsics. Efficiency Validation of 3G/4G WCDMA Air Interface Call Admission Control in OMNeT++ Environment. In *SoftCOM*, pages pp. 852–858, Split, Dubrovnik (Croatia), An-

- cona, Venice (Italy), October 2003. FESB-Split, ISBN 953-6114-64-X. [cf. COST-279, [TD\(03\)025](#)], [[173](#), [174](#)]. [21](#), [107](#), [146](#)
- [13] Nabil Benameur, Sara Oueslati, and James W. Roberts. Experimental Implementation of Implicit Admission Control. Technical Report 279TD(03)026, COST-279, 2003. [cf. COST-279, [TD\(03\)026](#)]. [21](#), [22](#)
 - [14] S. Ben Fredj, S. Oueslati-Boulahia, and J.W. Roberts. Measurement-based Admission Control for Elastic Traffic. In J. Moreira de Souza, N. L.S. da Fonseca, and E.A. de Souza e Silva, editors, *Teletraffic Engineering in the Internet Era*, pages 161–172. ITC 17, Elsevier, December 2001. [22](#)
 - [15] Eeva Nyberg and Samuli Aalto. How to Achieve Fair Differentiation. In *Proceedings of Networking 2002*, pages 1178–1183, Pisa, Italy, May 2002. IFIP-TC6, Springer-Verlag. [cf. COST-279, [TD\(01\)004](#)], [[175](#), [176](#)]. [22](#), [146](#)
 - [16] K. Kilkki and J. Ruutu. Simple Integrated Media Access - An Internet Service Based on Priorities . In *6th International Conference on Telecommunication Systems 1998*, 1998. [22](#)
 - [17] Andreas Terzis, J. Wang, J. Ogawa, and Lixia Zhang. A Two-Tier Resource Management Model for the Internet. In *Global Internet Symposium'99*, Dec. 1999. [23](#)
 - [18] Michael Menth, Stefan Kopf, Jens Milbrandt, and Joachim Charzinski. Introduction to Budget Based Network Admission Control Methods. In *28th Annual IEEE Conference on Local Computer Networks (LCN2003)*, Bonn, Germany, Oct. 2003. [23](#)
 - [19] Michael Menth, Sebastian Gehrsitz, and Jens Milbrandt. Fair Assignment of Efficient Network Admission Control Budgets. In *18th International Teletraffic Congress*, pages 1121–1130, Berlin, Germany, Sept. 2003. [23](#)
 - [20] Xipeng Xiao and Lionel M. Ni. Internet QoS: A Big Picture. *IEEE Network Magazine*, 13(2):8–18, March 1999. [24](#)
 - [21] A. Bak, W. Burakowski, F. Ricciato, S. Salsano, and H. Tarasiuk. Traffic handling in AQUILA QoS IP Network. In *21nd International Workshop on Quality of future Internet Services (QoFIS 2001)*, Coimbra, Portugal, Sept. 2001. [24](#)

- [22] Nick G. Duffield, Pawan Goyal, Albert G. Greenberg, Partho Pratim Mishra, K. K. Ramakrishnan, and Jacobus E. van der Merive. A Flexible Model for Resource Management in Virtual Private Networks. In *SIGCOMM*, pages 95 – 108, 1999. [24](#)
- [23] Michael Menth, Stefan Kopf, and Jens Milbrandt. A Performance Evaluation Framework for Network Admission Control Methods. In *IEEE Network Operations and Management Symposium (NOMS)*, Seoul, South Korea, April 2004. [24](#)
- [24] Michael Menth and Joachim Charzinski. Impact of Network Topology on the Performance of Network Admission Control Methods. In *International Workshop on Multimedia Interactive Protocols and Systems (MIPS2003)*, pages 195 – 206, Napoli, Italy, Nov. 2003. [24](#)
- [25] Michael Menth, Stefan Kopf, and Jens Milbrandt. Impact of Traffic Matrix and Routing on the Performance of Network Admission Control Methods. Technical Report, No. 307, University of Würzburg, Institute of Computer Science, Feb. 2003. [24](#)
- [26] G. A Politis, P Sampatakis, and I.S. Venieris. Design of a Multi-Layer Bandwidth Broker Architecture. In *Interworking*, Bergen, Norway, 2000. [25](#)
- [27] T. Engel, E. Nikolouzou, F. Ricciato, and P. Sampatakis. Analysis of Adaptive Resource Distribution Algorithm in the Framework of a Dynamic DiffServ IP Network. In *8th International Conference on Advances in Communications and Control (ComCon8)*, Crete, Greece, June 2001. [26](#)
- [28] Michael Menth. A Scalable Protocol Architecture for End-to-End Signaling and Resource Reservation in IP Networks. In *17th International Teletraffic Congress*, pages 211–222, Salvador da Bahia, Brazil, November 2001. [cf. COST-279, [TD\(01\)010](#)]. [26](#)
- [29] Michael Menth, Joachim Charzinski, and Stefan Kopf. Impact of Resilience Requirements on the Performance of Network Admission Control Methods. Technical Report 279TD(03)029, COST-279, 2003. [cf. COST-279, [TD\(03\)029](#)]. [26](#)
- [30] Michael Menth, Andreas Reifert, and Jens Milbrandt. Backup Capacity Minimization for Simple Protection Switching Mechanisms. Technical

- Report 279TD(03)046, COST-279, 2003. [cf. COST-279, [TD\(03\)046](#)]. [27](#), [42](#)
- [31] P. Olivier and N. Benameur. Flow Level IP Traffic Characterization. In *17th International Teletraffic Congress*, Salvador da Bahia, Brazil, December 2001. [cf. COST-279, [TD\(01\)014](#)]. [27](#), [62](#), [65](#)
- [32] J.W. Cohen. The multiple phase service network with generalized processor sharing. *Acta Informatica*, 12:245–284, 1979. [27](#), [66](#), [67](#), [91](#), [92](#)
- [33] L. Kleinrock. *Queueing Systems, Vol. 2: Computer applications*. J. Wiley & Sons, 1975. [27](#)
- [34] T. Bonald, P. Olivier, and J. Roberts. Dimensioning High Speed IP Access Networks. In *18th International Teletraffic Congress*, Berlin, Germany, Sept. 2003. [cf. COST-279, [TD\(03\)007](#)]. [28](#)
- [35] D. P. Heyman, T. V. Lakshman, and A. L. Neidhardt. A new method for analysing feedback-based protocols with applications to engineering Web traffic over the Internet. In *Proceedings of ACM Sigmetrics '97*, 1997. [28](#)
- [36] Hans van den Berg, Michel Mandjes, Remco van de Meent, Aiko Pras, Frank Roijers, and Pieter Venemans. QoS Aware Bandwidth Provisioning in IP Backbone Networks. Technical Report 279TD(03)034, COST-279, 2003. [cf. COST-279, [TD\(03\)034](#)]. [29](#), [51](#), [68](#)
- [37] S. Ben Fredj, T. Bonald, A. Proutière, G. Régnié, and J.W. Roberts. Statistical bandwidth sharing: a study of congestion at flow level. In *Proceedings of ACM SIGCOMM*, San Diego, CA, USA, Aug. 2001. [29](#)
- [38] R. Vranken, R.D. van der Mei, R.E. Kooij, and J.L. van den Berg. Flow-level performance models for the TCP with QoS differentiation. In *Proceedings of the International Seminar on Telecommunication Networks and Teletraffic Theory*, pages 78–87, St. Petersburg, Russia, 2002. [cf. COST-279, [TD\(02\)010](#)]. [30](#), [65](#), [66](#), [92](#)
- [39] R.D. van der Mei, J.L. van den Berg, R. Vranken, and B.M.M. Gijsen. Sojourn times in multi-server processor sharing systems with priorities. *Performance Evaluation*, 54:249–261, 2003. [30](#)
- [40] R. D. van der Mei, J. L. van den Berg, R. Vranken, and B. M. M. Gijsen. Analysis of a Flow Level Model for TCP Behavior in Case Of Priority

- Queueing. Technical Report 279TD(01)012, COST-279, 2001. [cf. COST-279, [TD\(01\)012](#)]. [31](#), [65](#), [66](#), [91](#), [92](#)
- [41] F. Delcoigne, A. Proutière, and G. Régnié. Modelling Integration of Streaming and Data Traffic. In *15th ITC Specialist Seminar on Internet Traffic Engineering and Traffic Management*, Würzburg, Germany, July 2002. [cf. COST-279, [TD\(02\)019](#)]. [32](#), [65](#)
- [42] A. De Vendictis and A. Baiocchi. Investigating TCP Single Source Behavior in Time-Varying Capacity Network Scenarios. In *18th International Teletraffic Congress*, Berlin, Germany, August/September 2003. [cf. COST-279, [TD\(03\)012](#)]. [33](#)
- [43] T. V. Lakshman and U. Madhow. The performance of TCP/IP for networks with high bandwidth-delay products and random loss. *IEEE/ACM Trans. on Networking*, 5(3):336–350, 1997. [34](#)
- [44] J. Padhye, V. Firoiu and D. F. Towsley, and J. F. Kurose. Modeling TCP Reno performance: a simple model and its empirical validation. *IEEE/ACM Trans. on Networking*, 8(2):133–145, 2000. [34](#)
- [45] Daniel Z. Lenardic, Branka Zovko-Cihlar, and Mislav Grgic. Analysis of Network Buffering Effects on TCP/IP Protocol Behavior. Technical Report 279TD(03)014, COST-279, 2003. [cf. COST-279, [TD\(03\)014](#)]. [34](#)
- [46] E. Plasser, T. Ziegler, and P. Reichl. On the Non Linearity of the RED Drop Function. In *ICCC*, Mubai, India, August 2002. [cf. COST-279, [TD\(02\)032](#)]. [35](#)
- [47] S. Floyd and V. Jacobson. Random Early Detection gateways for congestion avoidance. *IEEE/ACM Trans. on Networking*, 1(4):397–413, 1993. [35](#)
- [48] A. K. Jena and A. Popescu. Traffic Engineering for Internet Applications. In *Internet Performance and Control of Network Systems II*, volume 4523, pages 67–78, Denver, USA, August 2001. SPIE. [cf. COST-279, [TD\(01\)007](#)]. [36](#)
- [49] Francisco-Javier Ramón, José Enríquez, Jorge Andrés, and Antonio Molínes. Multipath Routing with Dynamic Variance. Technical Report 279TD(02)043, COST-279, 2002. [cf. COST-279, [TD\(02\)043](#)]. [37](#), [40](#)

- [50] Ivan Gojmerac, Thomas Ziegler, Fabio Ricciato, and Peter Reichl. Adaptive Multipath Routing for Dynamic Traffic Engineering. In *IEEE Globecom 2003*, San Francisco, USA, 2003. [cf. COST-279, [TD\(03\)028](#)]. 38, 40
- [51] Jens Milbrandt, Dirk Staehle, Stefan Köhler, and Les Berry. Decomposition of Large IP Networks for Routing Optimization. Technical Report 279TD(02)005, COST-279, 2002. [cf. COST-279, [TD\(02\)005](#)]. 39
- [52] Stefan Köhler and Andreas Binzenhöfer. MPLS Traffic Engineering in OSPF Networks — A Combined Approach. In *ITC 18*, Berlin, Germany, August/September 2003. [cf. COST-279, [TD\(03\)019](#)]. 40
- [53] Riikka Susitaival and Samuli Aalto. Providing Differentiated Services by Load Balancing and Scheduling in MPLS Networks. Technical Report 279TD(03)003, COST-279, 2003. [cf. COST-279, [TD\(03\)003](#)]. 40
- [54] R. Susitaival. Load Balancing by MPLS in Differentiated Services Networks. In *Art-QoS*, 2003. [cf. COST-279, [TD\(02\)041](#)]. 40
- [55] N. Akar, I. Hokelek, M. Atik, and E. Karasan. A reordering-free multipath traffic engineering architecture for diffserv/MPLS networks. In *IEEE Workshop on IP Operations and Management*, Kansas City, USA, October 2003. [cf. COST-279, [TD\(02\)026](#)]. 40
- [56] M. Nord, J. Soler Lucas, and V. Baek Iversen. Simulation and Performance Analysis of a GMPLS Lambda Scheduler. In *The Internet Protocol and Optical Networking Workshop*, Grasmere, UK, September 2002. [cf. COST-279, [TD\(03\)006](#)]. 42
- [57] M. Ajmone Marsan, M. Franceschinis, E. Leonardi, F. Neri, and A. Tarello. Instability Phenomena in Underloaded Packet Networks with QoS Schedulers. In *INFOCOM*, San Francisco, USA, March/April 2003. [cf. COST-279, [TD\(02\)027](#)]. 43
- [58] W. Burakowski and M. Fudala. Priority Forcing Scheme: A New Strategy for Getting Better than Best Effort Service in IP-based Network. In *Internet Technologies, Applications and Societal Impact*, Kluwer Academic Publishers, Wroclaw, Poland, October 2002. [cf. COST-279, [TD\(02\)039](#)]. 44
- [59] Rene Boel and Stijn De Vuyst. Prediction Based Resource Allocation, a Simulation Experiment. Technical Report 279TD(02)008, COST-279, 2002. [cf. COST-279, [TD\(02\)008](#)]. 44, 73

- [60] Oznur Ozkasap and Mine Caglar. Traffic Behavior of Scalable Multicast: Self-similarity and Protocol Dependence. In *18th International Teletraffic Congress*, Berlin, Germany, August/September 2003. [cf. COST-279, [TD\(02\)038](#)]. [45](#), [53](#), [54](#)
- [61] Oznur Ozkasap and Mine Caglar. Traffic Characterization of Scalable Multicasting in the case of a Self-Similar Source (Poster). In *ACM SIGCOMM*, Karlsruhe, Germany, August 2003. [cf. COST-279, [TD\(03\)036](#)]. [45](#), [53](#), [54](#)
- [62] Jouni Karvo and Samuli Aalto. Using Multicast or a Combination of Unicast and Broadcast for Transmitting Popular Content. Technical Report 279TD(03)020, COST-279, 2003. [cf. COST-279, [TD\(03\)020](#)]. [46](#)
- [63] Oznur Ozkasap. Scalability and Robustness of Pull-Based Anti-Entropy Distribution Model. In *ISCIS XVIII (18th International Symposium on Computer and Information Sciences)*, Antalya, Turkey, November 2003. [cf. COST-279, [TD\(03\)050](#)]. [46](#)
- [64] H. Reittu and I. Norros. On the Effect of Very Large Nodes in Internet Graphs. In *Globecom 2002*, Taipei, Taiwan, 2002. [cf. COST-279, [TD\(02\)002](#)]. [47](#)
- [65] H. Reittu and I. Norros. On the power law random graph model of massive data networks. *Performance Evaluation*, 2003. To appear. [cf. COST-279, [TD\(02\)020](#)]. [47](#)
- [66] I. Norros and H. Reittu. Architectural features of the power-law random graph model of Internet: Notes on soft hierarchy, vulnerability and multicasting. In *ITC-18*, Berlin, September 2003. [cf. COST-279, [TD\(03\)016](#)]. [48](#)
- [67] K. Salamatian, B. Baynat, and T. Bugnazet. Cross Traffic Estimation by Loss Process Analysis. In *15th ITC Specialist Seminar on Internet Traffic Engineering and Traffic Management*, Würzburg, Germany, July 2002. [cf. COST-279, [TD\(02\)015](#)]. [50](#), [51](#), [52](#), [53](#), [72](#)
- [68] Patrik Carlsson, Markus Fiedler, Kurt Tutschku, Stefan Chevul, and Arne A. Nilsson. Obtaining reliable bit rate measurements in SNMP-managed networks. In P. Tran-Gia and J. Roberts, editors, *Proceedings of 15th ITC Specialist Seminar on Internet Traffic Engineering and Traffic Management*, pages 114–123, Würzburg, Germany, July 2002. [cf. COST-279, [TD\(02\)021](#)]. [51](#)

- [69] J. Kilpi and I. Norros. Testing the Gaussian Character of Access Network Traffic. Technical Report 279TD(01)003, COST-279, 2001. [cf. COST-279, [TD\(01\)003](#)], [83]. [51](#), [56](#), [136](#)
- [70] A. Veres, Zs. Kenesi, S. Molnár, and G. Vattay. On the Propagation of Long Range Dependence in the Internet. In *ACM SIGCOM 2000*, Stockholm, Sweden, August/September 2000. [cf. COST-279, [TD\(01\)016](#)], [82]. [51](#), [55](#), [136](#)
- [71] S. Molnár and G. Terdik. A General Fractal Model of Internet Traffic. In *The 26th Annual IEEE Conference on Local Computer Networks*, Tampa, USA, November 2001. [cf. COST-279, [TD\(02\)004](#)]. [51](#), [52](#), [70](#)
- [72] Jorma Kilpi. A Portrait of a GPRS/GSM Session. Technical Report 279TD(02)040, COST-279, 2002. [cf. COST-279, [TD\(02\)040](#)], [84]. [51](#), [57](#), [101](#), [136](#)
- [73] Markus Fiedler, Kurt Tutschku, Patrik Carlsson, and Arne A. Nilsson. Identification of performance degradation in IP networks using throughput statistics. In J. Charzinski, R. Lehnert, and P. Tran Gia, editors, *Providing Quality of Service in Heterogeneous Environments. Proceedings of the 18th International Teletraffic Congress (ITC-18)*, pages 399–407, Berlin, Germany, September 2003. [cf. COST-279, [TD\(03\)021](#)]. [51](#), [68](#)
- [74] Jorma Kilpi. Distributional Properties of GPRS/GSM Session Volumes and Durations. Technical Report 279TD(03)023, COST-279, 2003. [cf. COST-279, [TD\(03\)023](#)]. [51](#), [57](#), [101](#)
- [75] Michela Meo, Marco Ajmone Marsan, Luca Muscariello, Marco Mellia, and Renato Lo Cigno. A Simple Markovian Approach to Model Internet Traffic at Edge Routers. Technical Report 279TD(03)032, COST-279, 2003. [cf. COST-279, [TD\(03\)032](#)]. [51](#), [64](#)
- [76] Kurt Tutschku and Phuoc Tran-Gia. A Traffic Profile of the eDonkey Filesharing Service. Technical Report 279TD(03)049, COST-279, 2003. [cf. COST-279, [TD\(03\)049](#)]. [51](#), [59](#), [60](#)
- [77] Meta Search Inc. eDonkey2000 Home Page. <http://www.edonkey2000.com/>. [51](#), [58](#), [60](#)

- [78] Patrik Carlsson, Markus Fiedler, and Anders Ekberg. On an Implementation of a Distributed Passive Measurement Infrastructure. Technical Report 279TD(03)042, COST-279, 2003. [cf. COST-279, [TD\(03\)042](#)]. [51](#)
- [79] Kurt Tutschku and Hermann deMeer. A measurement study on signaling on gnutella overlay networks. In *Fachtagung - Kommunikation in Verteilten Systemen (KiVS) 2003*, pages 295–306, Leipzig, Germany, 2 2003. [cf. COST-279, [TD\(03\)004](#)]. [52](#), [59](#)
- [80] J. Bachmann, M. Matthes, O. Drobnik, and U. R. Krieger. Mobility and QoS-Management for Adaptive Applications. In *11th International World Wide Web Conference WWW2002*, Honolulu, USA, may 2002. [cf. COST-279, [TD\(02\)001](#)]. [53](#), [58](#), [116](#)
- [81] Guoqiang Hu, Klaus Dolzer, and Christoph Gauger. Does Burst Assembly Really Reduce the Self-Similarity? In *Optical Fiber Communication 2003*, pages 124–126, Atlanta, Georgia, USA, March 2003. [cf. COST-279, [TD\(03\)030](#)]. [55](#)
- [82] A. Veres, Zs. Kenesi, S. Molnár, and G. Vattay. Tcp’s role in the propagation of self-similarity in the internet. *Computer Communications, Special Issue on Performance Evaluation of IP Networks*, 26(8):899–913, May 2003. [cf. COST-279, [TD\(01\)016](#)], [70]. [55](#), [135](#)
- [83] J. Kilpi and I. Norros. Testing the Gaussian Approximation of Aggregate Traffic. In *The 2nd Internet Measurement Workshop*, Marseille, France, 2002. [cf. COST-279, [TD\(01\)003](#)], [69]. [56](#), [135](#)
- [84] Jorma Kilpi. A Portrait of a GPRS/GSM Session. In *18th International Teletraffic Congress*, Berlin, Germany, August/September 2003. [cf. COST-279, [TD\(02\)040](#)], [72]. [57](#), [101](#), [135](#)
- [85] Anonymous. The gnutella protocol specification v0.4. available at <http://dss.clip2.com>, Clip2 Distributed Search Solutions, 2001. [58](#)
- [86] Sharman Networks. Kazaa media desktop – <http://www.kazaa.com/>. [58](#)
- [87] N. Markovitch and U. R. Krieger. The Estimation of Heavy-Tailed Probability Density Functions, Their Mixtures and Quantiles. *Computer Networks*, 2002. [cf. COST-279, [TD\(01\)001](#)]. [61](#)
- [88] Udo Krieger and N. M. Markovitch. On-Line Estimation of Heavy-Tailed Traffic Characteristics in Web Data Mining. Technical Report 279TD(03)039, COST-279, 2003. [cf. COST-279, [TD\(03\)039](#)]. [61](#)

- [89] Guy Latouche and Marie-Ange Remiche. A MAP-based Poisson Cluster Model for Web Traffic. *Performance Evaluation*, 49:359–370, 2002. [cf. COST-279, [TD\(02\)016](#)]. [62](#)
- [90] Z. Liu, N. Niclausse, and C. Jalpa-Villanueva. Traffic model and performance evaluation of web servers. *Performance Evaluation*, 46:77–100, 2001. [62](#)
- [91] G. Latouche, M.-A. Remiche, and P. Taylor. Transient markov arrival processes. *The Annals of Applied Probability*, 13:628–640, 2003. [62](#)
- [92] A. Andersen and B. Nielsen. A markovian approach for modeling packet traffic with long-range dependence. *Journal on Selected Areas in Communications*, 16(5):719–732, 1998. [62](#)
- [93] Aleš Švigelj, Mihael Mohorčič, and Gorazd Kandus. Traffic Class Dependent Routing in Packet-Switched Non-Geostationary ISL Networks. In *PIMRC 2002: IEEE International Symposium on Personal, Indoor and Mobile Radio Communications*, volume 3, pages 1382–1386, Lisbon, Portugal, September 2002. [cf. COST-279, [TD\(03\)024](#)], [[177](#), [178](#), [179](#), [180](#), [181](#)]. [63](#), [119](#), [120](#), [146](#), [147](#)
- [94] Paulo Salvador, Antonio Pacheco, and Rui Valadas. Multiscale Fitting Procedure using Markov Modulated Poisson Processes. *Telecommunication Systems Journal*, 23(1-2):123–148, June 2003. [cf. COST-279, [TD\(02\)018](#)]. [63](#)
- [95] V. Inghelbrecht, B. Steyaert, S. Wittevrongel, and H. Bruneel. Analytic study of the interdeparture time characteristics in a multistage network. In *ITC 18*, volume 5b, pages 1191–1200, Berlin, Germany, August/September 2003. [cf. COST-279, [TD\(02\)030](#)]. [64](#), [97](#)
- [96] A. Feldmann. Characteristics of tcp connection arrivals. In K. Park and W. Willinger, editors, *Self-similar network traffic and performance evaluation*. J. Wiley & Sons, 2000. [65](#)
- [97] Hans van den Berg, Robert J. Kooij, Michel Mandjes, Pasi Lassila, and Robert van der Mei. An integrated packet/flow model for TCP performance analysis. In *ITC 18*, pages 651–660, Berlin, Germany, August/September 2003. [cf. COST-279, [TD\(02\)034](#)]. [67](#)
- [98] J. Padhye, V. Firoiu, D. Towsley, and J. Kurose. Modeling tcp throughput: a simple model and its empirical validation. In *Proceedings of ACM SIGCOMM'98*, Vancouver, CA, September 1998. [67](#)

- [99] Remco Litjens, Hans van den Berg, Richard J. Boucherie, Frank Roijers, and Maria Fleuren. Performance Analysis of Wireless LANs: an Integrated Packet/Flow Level Approach. Technical Report 279TD(03)001, COST-279, 2003. [cf. COST-279, [TD\(03\)001](#)]. [67](#), [115](#)
- [100] Kathleen Spaey, Tom Hofkens, and Chris Blondia. Timescales in Models for Bursty Traffic. Technical Report 279TD(03)002, COST-279, 2003. [cf. COST-279, [TD\(03\)002](#)], [[101](#)]. [67](#), [68](#), [79](#), [138](#)
- [101] Kathleen Spaey, Danny De Vleeschauwer, Tom Hofkens, and Chris Blondia. Timescales in Models for Bursty Traffic. In Mohammad S. Obaidat, Franco Davoli, Erina Ferro, and Ibrahim Onyuksel, editors, *2003 International Symposium on Performance Evaluation of Computer and Telecommunication Systems (SPECTS 2003)*, volume 35(4) of *Simulation Series*, pages 465–473. SCS, July 2004. [cf. COST-279, [TD\(03\)002](#)], [[100](#)]. [68](#), [138](#)
- [102] Markus Fiedler and Kurt Tutschku. Application of the Stochastic Fluid Flow Model for bottleneck identification and classification. In *Proceedings of 2003 Design, Analysis, and Simulation of Distributed Systems (DASD 2003)*, pages 35–42, Orlando, USA, April 2003. [cf. COST-279, [TD\(02\)046](#)]. [68](#), [69](#), [87](#)
- [103] Markus Fiedler, Patrik Carlsson, and Arne A. Nilsson. Voice and multifractal data traffic in the Internet. In *Proceedings of the 26th Annual IEEE Conference on Local Computer Networks (LCN 2001)*, pages 426–431, Tampa, USA, November 2001. [cf. COST-279, [TD\(01\)008](#)]. [69](#), [84](#)
- [104] P. Mannersalo, I. Norros, and R. Riedi. Multifractal products of stochastic processes: A preview. Technical Report 257TD(99)31, COST-257, 1999. [69](#)
- [105] Paulo Salvador and Rui Valadas António Nogueira. Joint Characterization of the Packet Arrival and Packet Size Processes of Multifractal Traffic based on Stochastic L-Systems. In *International Teletraffic Congress*, Berlin, Germany, September 2003. [cf. COST-279, [TD\(03\)008](#)]. [69](#), [70](#)
- [106] A. Medina, N. Taft, K. Salamatian, S. Bhattacharyya, and C. Diot. Traffic matrix estimation: Existing techniques and new directions. In

- ACM SIGCOMM 2002*, Pittsburg, USA, August 2002. [cf. COST-279, [TD\(02\)014](#)]. 71
- [107] P. Mannersalo. Some Notes on Prediction of Teletraffic. In *15th ITC Specialist Seminar*, pages 220–229, Würzburg, Germany, 2002. [cf. COST-279, [TD\(02\)025](#)]. 72
- [108] I. Norros. Most Probable Path Techniques for Gaussian Queueing Systems. In *Networking 2002*, Pisa, Italy, 2002. [cf. COST-279, [TD\(02\)003](#)]. 73, 88
- [109] Joris Walraevens, Bart Steyaert, and Herwig Bruneel. A Single-Server Queue with a Priority Scheduling Discipline: Performance Study. Technical Report 279TD(01)006, COST-279, 2001. [cf. COST-279, [TD\(01\)006](#)]. 76
- [110] J. Walraevens, B. Steyaert, and H. Bruneel. Performance analysis of a GI-G-1 preemptive resume priority buffer. In *Proceedings of the Networking 2002 Conference, Pisa, May 19-24*, LNCS 2345, pages 745–756, 2002. 77
- [111] J. Walraevens, B. Steyaert, and H. Bruneel. Analysis of a preemptive repeat priority buffer with resampling. In *Proceedings of the International Network Optimization Conference (INOC 2003)*, Evry, October 27-29, pages 581–586, 2003. 77
- [112] J. Walraevens, B. Steyaert, and H. Bruneel. Delay characteristics in discrete-time GI-G-1 queues with non-preemptive priority queueing discipline. *Performance Evaluation*, 50(1):53–75, 2002. 77
- [113] D. Fiems, B. Steyaert, and H. Bruneel. Discrete-Time Queues with General Service Times and General Server Interruptions. In *SPIE*, volume 4211, pages 93–104, Boston, USA, 2000. [cf. COST-279, [TD\(01\)005](#)]. 77
- [114] D. Fiems, S. de Vuyst, and H. Bruneel. The Combined Gated-Exhaustive Vacation System in Discrete Time. *Performance Evaluation*, pages 227–239, 2002. [cf. COST-279, [TD\(03\)017](#)]. 78
- [115] P. Gao, S. Wittevrongel, and H. Bruneel. Analysis of Discrete-Time Buffers with Geometric Service Times and Multiple Servers. In *High Performance Computing Symposium*, pages 294–299, San Diego, April 2002. [cf. COST-279, [TD\(02\)035](#)], [182]. 78, 147

- [116] P. Gao, S. Wittevrongel, and H. Bruneel. Discrete-Time Multi-server Buffer Systems with Correlated Arrivals and Geometric Service Times. In *the Conference on Design, Analysis, and Simulation of Distributed Systems*, pages 27–34, Orlando, April 2003. [cf. COST-279, [TD\(03\)041](#)]. [78](#)
- [117] P. Gao, S. Wittevrongel, and H. Bruneel. Delay against system contents in discrete-time G/Geom/c queue. *Electronics Letters*, 39(17):1290–1292, 2003. [79](#)
- [118] K. Laevens and H. Bruneel. Analysis of a Single-Wavelength Optical Buffer. In *INFOCOM 2003*, San Francisco, USA, April 2003. [cf. COST-279, [TD\(03\)015](#)]. [80](#)
- [119] J. Walraevens, S. Wittevrongel, and H. Bruneel. An Analytic Technique to Evaluate the Performance of Optical Packet Switches. In *7th IFIP Working Conference on Optical Network Design & Modelling*, volume 2, pages 1171–1185, Budapest, Hungary, February 2003. [cf. COST-279, [TD\(02\)036](#)], [[183](#)]. [81](#), [147](#)
- [120] Veronique Inghelbrecht, Bart Steyaert, and Herwig Bruneel. Study of the burstification mechanism of an OBS edge router. In *7th IFIP Working Conference on Optical Network Design & Modelling Conference (ONDM2003)*, volume 2, pages 1221–1239, Budapest, Hungary, February 2003. [cf. COST-279, [TD\(03\)037](#)]. [81](#)
- [121] Krzysztof Tworus, Stijn de Vuyst, Sabine Wittevrongel, and Herwig Bruneel. Queueing Analysis of the Stop-and-Wait ARQ Protocol in a Wireless Environment. Technical Report 279TD(03)043, COST-279, 2003. [cf. COST-279, [TD\(03\)043](#)]. [82](#)
- [122] V. Ramaswami. Matrix analytic methods for stochastic fluid flows. In D. Smith and P. Hey, editors, *Proceedings of the 16th International Teletraffic Congress, Edinburgh, UK*, pages 1019–1030. Elsevier Science B.V., 1999. [83](#)
- [123] G. Latouche and V. Ramaswami. A logarithmic reduction algorithm for quasi-birth-and-death processes. *J. Appl. Probab.*, 30:650–674, 1993. [83](#)
- [124] A. da Silva Soares and G. Latouche. Further Results on the Similarity Between Fluid Queues and QBDs. In G. Latouche and P. Taylor, editors,

- Matrix-Analytic Methods Theory and Applications—Proceedings of the 4th International Conference on Matrix-Analytic Methods*, pages 89–106. World Scientific, 2002. [cf. COST-279, [TD\(01\)002](#)]. 83
- [125] D. Anick, D. Mitra, and M.M. Sondhi. Stochastic theory of a data-handling system with multiple sources. *The Bell System Technical Journal*, 61(8):1871–1894, 1982. 83, 87
- [126] N. Akar and K. Sohraby. Numerically Stable Solution of Large-Scale Finite Fluid Queues. In *International Teletraffic Congress*, Berlin, Germany, September 2003. [cf. COST-279, [TD\(03\)011](#)],[184]. 83, 84, 147
- [127] Daniel Zaragoza and Carlos A. Carvalho Belo. Queuing Behavior of Multiplexed General ON-OFF Sources: An Engineering Perspective. Technical Report 279TD(02)007, COST-279, 2002. [cf. COST-279, [TD\(02\)007](#)]. 84, 85
- [128] Michel Mandjes, Debasis Mitra, and Werner Scheinhardt. Models of network access using feedback fluid queues. *Queueing Systems*, 44:365–398, 2003. [cf. COST-279, [TD\(02\)011](#)], [185, 186]. 85, 147
- [129] Detlef Sass, Klaus Dolzer, Stefan Bodamer, and Martin Lorang. Analytic Fluid Flow Approach for Fair Queueing Systems. Technical Report 279TD(03)033, COST-279, 2003. [cf. COST-279, [TD\(03\)033](#)]. 86
- [130] P. Mannersalo and I. Norros. A Most Probable Path Approach to Queueing Systems with General Gaussian Input. *Computer Networks*, 40:399–412, 2002. [cf. COST-279, [TD\(01\)013](#)]. 88
- [131] Michel Mandjes, Petteri Mannersalo, Ilkka Norros, and Miranda van Uitert. Most Probable Busy Period Paths in Gaussian Queues. Technical Report 279TD(03)048, COST-279, 2003. [cf. COST-279, [TD\(03\)048](#)]. 89
- [132] Gerhard Haßlinger and Markus Fiedler. Network dimensioning for Gaussian traffic aggregated from Markovian on-off sources. In *SPIE ITCOM 2002*, Boston, USA, July/August 2002. [cf. COST-279, [TD\(02\)013](#)]. 90
- [133] T. Bonald en J.W. Roberts. Performance of bandwidth sharing mechanisms for service differentiation in the Internet. In *Proceedings ITC Specialists Seminar on IP Measurement, Modeling and Management*, Monterey, pages 22–1 – 22–10, 2000. 91

- [134] J.V.L. Beckers, I. Hendrawan, R.E. Kooij, and R.D. van der Mei. Generalized processor sharing models for Internet access lines. In *Proceedings 9th IFIP conference on Performance Modeling and Evaluation of ATM & IP Networks, Budapest*, pages 101–112, 2001. 91
- [135] A. Riedl, T. Bauschert, M. Perske, and A. Probst. Investigation of the M/G/R processor sharing system for dimensioning IP access networks with elastic traffic. In *Proceedings 1st Polish-German Symposium of Telecommunication Systems*, 2000. 91
- [136] S.F. Yashkov. Mathematical problems in the theory of processor-sharing queueing systems. *Journal of Soviet Mathematics*, 58:101–147, 1992. 91
- [137] S.F. Yashkov. Processor-sharing queues: some progress in analysis. *Queueing Systems*, 2:1–17, 1987. 91
- [138] Remco Litjens M. Sc., Hans van den Berg, and Richard J. Boucherie. Throughput Measures for Processor Sharing Models. Technical Report 279TD(03)022, COST-279, 2003. [cf. COST-279, TD(03)022]. 92, 93
- [139] C. Douligeris. Multiobjective flow control in telecommunication networks. In *Proceedings of INFOCOM '92, Florence, Italy*, 1992. 92
- [140] A.A. Kherani and A. Kumar. Performance analysis of TCP with nonpersistent sessions. In *Proceedings of the Workshop on Modelling of flow and congestion control, Paris, France*, 2000. 92
- [141] A.A. Kherani and A. Kumar. Stochastic models for throughput analysis of randomly arriving elastic flows in the Internet. In *Proceedings of INFOCOM '02, New York, USA*, 2002. 92
- [142] S. Ben Fredj, T. Bonald, A. Proutiere, G. Régnié, and J.W. Roberts. Statistical bandwidth sharing: a study of congestion at flow level. In *Proceedings of SIGCOMM '01, San Diego, USA*, 2001. 93
- [143] N. Benameur, S. Ben Fredj, F. Delcoigne, S. Oueslati-Boulaiahia, and J.W. Roberts. Integrated admission control for streaming and elastic traffic. In *Proceedings of the 2nd International workshop on Quality of future Internet services, Coimbra, Portugal*, 2001. 93
- [144] T. Bonald and L. Massoulié. Impact of fairness on Internet performance. In *Proceedings of SIGMETRICS '01, Cambridge, USA*, 2001. 93

- [145] F. Delcoigne, A. Proutière, and G. Régnié. Modelling integration of streaming and data traffic. In *Proceedings of the ITC specialist seminar on Internet traffic engineering and traffic management*, Würzburg, Germany, 2002. 93
- [146] E. Kuumola, J. Resing, and J. Virtamo. Joint Distribution of Instantaneous and Averaged Queue Length in an M/M/1/K System. In Phuoc Tran-Gia and J. Roberts, editors, *15th ITC Specialist Seminar on Internet Traffic Engineering and Traffic Management*, pages 58–67, Würzburg, Germany, July 2002. [cf. COST-279, TD(02)022]. 93
- [147] Udo Krieger. The BMAP/G/1 Queue with Feedback Operating in Synchronous Random Environment as a Model for a Telecommunication Channel with Performance Fluctuation. Technical Report 279TD(03)038, COST-279, 2003. [cf. COST-279, TD(03)038]. 95, 96
- [148] Olav Øserbø Telenor FoU. An Approximative Method to Calculate the Distribution of End-to-End Delay in Packet Networks. Technical Report 279TD(02)033, COST-279, 2002. [cf. COST-279, TD(02)033]. 97
- [149] D. De Vleeschauwer, G.H. Petit, B. Steyaert, S. Wittevrongel, and H. Bruneel. Calculation of end-to-end delay quantile in network of M/G/1 queues. *Electronics Letters*, 37:535–536, 2001. 97
- [150] 3GPP. The third generation partnership project, 1998. 100
- [151] M. Meo, M. Ajmone Marsan, and C. Batetta. Resource Management Policies in GPRS Wireless Internet Access Systems. In *IEEE International Performance and Dependability Symposium*, Washington, USA, June 2002. [cf. COST-279, TD(02)028], [187]. 102, 147
- [152] Richard J. Boucherie and Aljaz Ule. Adaptive Dynamic Channel Borrowing in Road-Covering Mobile Networks. Technical Report 279TD(02)012, COST-279, 2002. [cf. COST-279, TD(02)012]. 102
- [153] Johan van Leeuwen, Samuli Aalto, and Jorma Virtamo. Load Balancing in Cellular Networks Using First Policy Iteration. Technical Report 279TD(02)023, COST-279, 2002. [cf. COST-279, TD(02)023]. 103
- [154] Dirk Staehle and Andreas Mäder. An Analytic Approximation of the Uplink Capacity in a UMTS Network with Heterogeneous Traffic. In *ITC-18*, Berlin, September 2003. [cf. COST-279, TD(03)018]. 106

- [155] Hans van den Berg, Aloysius Irwan Endrayanto, and Richard J. Boucherie. Characterizing CDMA Downlink Feasibility via Effective Interference. Technical Report 279TD(03)047, COST-279, 2003. [cf. COST-279, [TD\(03\)047](#)]. 107
- [156] Branka Zovko-Cihlar, Winton Afric, and Sonja Grgic. Interference in Direct Spread Sequence Mobile Communication System. Technical Report 279TD(03)013, COST-279, 2003. [cf. COST-279, [TD\(03\)013](#)]. 108
- [157] Dirk Staehle, Kenji Leibnitz, Klaus Heck, B. Schröder, A. Weller, and Phuoc Tran-Gia. Analytical Characterization of the Soft Handover Gain in UMTS. In *The IEEE Vehicular Technology Conference*, Atlantic City, USA, October 2001. [cf. COST-279, [TD\(02\)006](#)]. 109
- [158] Klaus Heck, Dirk Staehle, and Kenji Leibnitz. Diversity Effects on the Soft Handover Gain in UMTS Networks. In *IEEE Vehicular Technology Conference*, Vancouver, Canada, 2002. [cf. COST-279, [TD\(02\)042](#)]. 109
- [159] R. Litjens. The Impact of Mobility on UMTS Network Planning. *Computer Networks*, 38:497–515, 2002. [cf. COST-279, [TD\(01\)011](#)]. 110, 111
- [160] Sandor Szabo, Sándor Imre, and Alexandrosz Burulitisz. On the Accuracy of Mobility Modelling in Wireless Networks. In *JOHN VON NEUMANN PHD CONFERENCE*, 2003. [cf. COST-279, [TD\(03\)031](#)]. 111
- [161] Remco Litjens M. Sc. and Richard J. Boucherie. Performance Analysis of Downlink Shared Channels in a UMTS Network. Technical Report 279TD(02)045, COST-279, 2002. [cf. COST-279, [TD\(02\)045](#)]. 113
- [162] Remco Litjens M. Sc. and Hans van den Berg. Fair Adaptive Sheduling in Integrated Services UMTS Networks. Technical Report 279TD(02)016, COST-279, 2002. [cf. COST-279, [TD\(02\)016](#)]. 113
- [163] IEEE. ANSI/IEEE Std 802.11, 1999 Edition, 1999. 114
- [164] C. Blondia, O. Casals, Ll. Cerdà, N. Van den Wijngaert, G. Willems, and P. De Cleyn. Low Latency Handoff Mechanism and Their Implementation in an IEEE 802.11 Network. In *18th ITC*, Berlin, Germany, September 2003. [cf. COST-279, [TD\(03\)004](#)]. 115

- [165] C. Blondia, O. Casals, P. De Cleyn, and G. Willems. Performance analysis of IP Micro-Mobility Handoff Protocols. In *Proceedings of Protocols for High Speed Networks 2002 (PfHSN 2002)*, pages 211–226, Berlin, 2002. [115](#)
- [166] C. Blondia, O. Casals, Ll. Cerdà, and G. Willems. Performance Analysis of a Forwarding Scheme for Handoff in HAWAII. In *Proceedings of Networking 2002*, pages 504–514, Pisa, 2002. [115](#)
- [167] C. Blondia, O. Casals, N. Van den Wijngaert, and G. Willems. Performance analysis of smooth handoff in Mobile IP. In *Proceedings of MSWiM'2002*, Atlanta, USA, September 2002. [115](#)
- [168] O. Casals, Ll. Cerdà, G. Willems, C. Blondia, and N. Van den Wijngaert. Performance evaluation of the Post-Registration method, a low latency handoff in MIPv4. In *Proceedings of ICC*, 2003. [115](#)
- [169] C. Blondia, O. Casals, Ll. Cerdà, N. Van den Wijngaert, G. Willems, and P. de Cleyn. Performance Comparison of Low Latency Mobile IP Schemes. In *Proceedings of WiOpt'03*, pages 115–124, INRIA Sophia Antipolis, France, March 2003. [115](#)
- [170] R. Schulcz, S. Szabó, S. Imre, and L. Pap. The Effect of Radio Cell Size In wMATM Based Third Generation Mobile Systems. In *Services & Applications in the Wireless Public Infrastructure 2001*, pages 59–72, Paris, France, July 2001. HERMES Scientific Publications. [cf. COST-279, [TD\(01\)015](#)]. [118](#)
- [171] M. Szalay and S. Imre. Reliability Modelling of Tree Topology IP Micro Mobility Networks. In *EUNICE 2002 7th Open European Summer School on Adaptable Networks and Teleservices*, pages 63–69, Trondheim, Norway, September 2002. Tapir Uttrykk. [cf. COST-279, [TD\(03\)009](#)], [[188](#)]. [119](#), [147](#)
- [172] H. T. Tran and T. Ziegler. An Admission Control Scheme for Voice Traffic over IP Networks. In *LNCS 2720, Proceedings of the 6th IEEE International Conference on High Speed Networks and Multimedia Communications HSNMC'03*, pages 353–364, July 2003. [cf. COST-279, [TD\(02\)037](#)], [[7](#)]. [128](#)
- [173] S. Imre. Dynamically optimised chernoff bound based cac for 3g/4g wcdma systems. In *Microcoll*, pages pp. 27–30, Budapest, Hungary,

- September 2003. Microcoll, ISBN 963-212-166-X. [cf. COST-279, [TD\(03\)025](#)], [[12](#), [174](#)]. [129](#), [146](#)
- [174] S. Imre, K. Hank, P. Petrs, and R. Tancsics. Efficient call admission control method for 3g/4g wcdma networks. In *CONTEL*, pages pp. 293–300, Zagreb, Croatia, June 2003. CIP Zagreb, ISBN 953-184-055-5. [cf. COST-279, [TD\(03\)025](#)], [[12](#), [173](#)]. [129](#), [146](#)
- [175] Samuli Aalto and Eeva Nyberg. Flow level models of DiffServ packet level mechanisms. In *Proceedings of the Sixteenth Nordic Teletraffic Seminar, NTS 16*, pages 194–205, Espoo, Finland, August 2002. [cf. COST-279, [TD\(01\)004](#)], [[15](#), [176](#)]. [129](#), [146](#)
- [176] E. Nyberg and S. Aalto. Differentiation and interaction of traffic: a flow level study. In *Proceedings of International Workshop, Art-Qos 2003*, pages 276–290, Warsaw, Poland, March 2003. [cf. COST-279, [TD\(01\)004](#)], [[15](#), [175](#)]. [129](#), [146](#)
- [177] Aleš Švigelj, Mihael Mohorčič, and Gorazd Kandus. Adaptive Packet Routing Based on Traffic Class Differentiation in Intersatellite Link Networks. *WSEAS Transactions on Communications*, 1:138–143, 2002. [cf. COST-279, [TD\(03\)024](#)], [[93](#), [178](#), [179](#), [180](#), [181](#)]. [137](#), [146](#), [147](#)
- [178] Aleš Švigelj, Mihael Mohorčič, and Gorazd Kandus. Traffic Class Dependent Routing in Packet-Switched Non-Geostationary ISL Networks. In E. Del Re, editor, *Mobile and personal satellite communications 5: Proceedings of the Fifth European Workshop on Mobile/Personal Satcoms (EMPS 2002)*, pages 45–52, Baveno, Italy, September 2002. [cf. COST-279, [TD\(03\)024](#)], [[93](#), [177](#), [179](#), [180](#), [181](#)]. [137](#), [146](#), [147](#)
- [179] Aleš Švigelj, Mihael Mohorčič, and Gorazd Kandus. Traffic Class Dependent Routing in ISL Network with Adaptive Forwarding Based on Local Link Load Information. In *Satellite communications - From fade mitigation to service provision: International Workshop of COST Actions 272 and 280*, pages 395–402, Noordwijk, The Netherlands, May 2003. COST-272 / COST-280, ESTEC/ESA. [cf. COST-279, [TD\(03\)024](#)], [[93](#), [177](#), [178](#), [180](#), [181](#)]. [137](#), [146](#), [147](#)
- [180] Gorazd Kandus, Aleš Švigelj, and Mihael Mohorčič. The Impact of Different Scheduling Policies on Traffic Class Dependent Routing in Intersatellite Link Network. In *International Conference on Advanced*

- Satellite Mobile Systems (ASMS 2003)*, ESA SP-541, Frascati, Italy, July 2003. ESA/ESTEC. [cf. COST-279, [TD\(03\)024](#)], [[93](#), [177](#), [178](#), [179](#), [181](#)]. [137](#), [146](#), [147](#)
- [181] Aleš Švigelj, Mihael Mohorčič, Gorazd Kandus, Andrej Kos, Matevž Pustišek, and Janez Bešter. Routing in ISL networks Considering Empirical IP Traffic. *IEEE Journal on Selected Areas in Communications*, to appear in February 2004. [cf. COST-279, [TD\(03\)024](#)], [[93](#), [177](#), [178](#), [179](#), [180](#)]. [137](#), [146](#), [147](#)
- [182] P. Gao, S. Wittevrongel, and H. Bruneel. Discrete-Time Multiserver Queues with Geometric Service Times. *Computers & Operations Research*, 31(1):81–99, 2004. [cf. COST-279, [TD\(02\)035](#)], [[115](#)]. [139](#)
- [183] J. Walraevens, S. Wittevrongel, and H. Bruneel. Calculation of the Packet Loss in Optical Packet Switches: an Analytic Technique. *AEU (International Journal of Electronics and Communications)*, 57(4):270–276, 2003. [cf. COST-279, [TD\(02\)036](#)], [[119](#)]. [140](#)
- [184] N. Akar and K. Sohraby. Infinite/finite Markov fluid queues: A unified analysis. *Journal of Applied Probability*, 2004. to appear, [cf. COST-279, [TD\(03\)011](#)], [[126](#)]. [141](#)
- [185] Michel Mandjes, Debasis Mitra, and Werner Scheinhardt. A simple model of network access: feedback adaptation of rates and admission control. *Computer Networks*, 41:489–504, 2003. [cf. COST-279, [TD\(02\)011](#)], [[128](#), [186](#)]. [141](#), [147](#)
- [186] Michel Mandjes, Debasis Mitra, and Werner Scheinhardt. A simple model of network access: feedback adaptation of rates and admission control. In *Proceedings INFOCOM 2002*, pages 3–12, New York, US, July 2002. [cf. COST-279, [TD\(02\)011](#)], [[128](#), [185](#)]. [141](#), [147](#)
- [187] M. Meo and M. Ajmone Marsan. Resource Management Policies in GPRS Systems. *Performance Evaluation*, 2002. [cf. COST-279, [TD\(02\)028](#)], [[151](#)]. [143](#)
- [188] M. Szalay and S. Imre. Reliability Considerations of IP Micro Mobility Networks. In *Design of Reliable Communication Networks DRCN*, pages 72–77, Budapest, Hungary, October 2001. [cf. COST-279, [TD\(03\)009](#)], [[171](#)]. [145](#)

- [189] E. Karasan, O. Karasan, N. Akar, and M. Pinar. Topology Design in Virtual Private Networks. In *INFORMS National Meeting*, San Jose, USA, November 2002. [cf. COST-279, [TD\(03\)010](#)], [[190](#)]. [148](#)
- [190] E. Karasan, O. Karasan, N. Akar, and M. Pinar. Mesh Topology Design in Overlay Virtual Private Networks. *Electronics Letters*, 38(16):939–941, 2002. [cf. COST-279, [TD\(03\)010](#)], [[189](#)]. [148](#)

List of Temporary Documents

- [[TD\(01\)001](#)] Udo Krieger and Natalia M. Markovich. The Estimation of Heavy-Tailed Probability Density Functions. Technical Report 279TD(01)001, COST-279, 2001.
- [[TD\(01\)002](#)] Ana de Silva Soares and Guy Latouche. Algorithmic Approach to Fluid Queues. Technical Report 279TD(01)002, COST-279, 2001.
- [[TD\(01\)003](#)] J. Kilpi and I. Norros. Testing the Gaussian Character of Access Network Traffic. Technical Report 279TD(01)003, COST-279, 2001.
- [[TD\(01\)004](#)] Eeva Nyberg, Jorma Virtamo, and Samuli Aalto. Relating Flow Level Requirements to DiffServ Packet Level Mechanisms. Technical Report 279TD(01)004, COST-279, 2001.
- [[TD\(01\)005](#)] Dieter Fiems, Herwig Bruneel, and Bart Steyaert. Analysis of a Discrete-Time Queueing Model with Server Interruptions Modeling Preemptive Priority Systems. Technical Report 279TD(01)005, COST-279, 2001.
- [[TD\(01\)006](#)] Joris Walraevens, Bart Steyaert, and Herwig Bruneel. A Single-Server Queue with a Priority Scheduling Discipline: Performance Study. Technical Report 279TD(01)006, COST-279, 2001.
- [[TD\(01\)007](#)] A. K. Jena and A. Popescu. Traffic Engineering for Internet Services. Technical Report 279TD(01)007, COST-279, 2001.
- [[TD\(01\)008](#)] Markus Fiedler, Patrik Carlsson, and Arne A. Nilsson. Fluid flow analysis for voice over IP and multi-fractal data traffic. Technical Report 279TD(01)008, COST-279, 2001.

- [TD(01)009] Klaus Dolzer and Christoph Gauger. On Burst Assembly in Optical Burst Switching Networks - A Performance Evaluation of Just Enough-Time. Technical Report 279TD(01)009, COST-279, 2001.
- [TD(01)010] Michael Menth. A Scalable Protocol Architecture for End-to-End Signaling and Resource Reservation in IP Networks. Technical Report 279TD(01)010, COST-279, 2001.
- [TD(01)011] Remco Litjens M. Sc. The impact of Mobility of UMTS Network Planning. Technical Report 279TD(01)011, COST-279, 2001.
- [TD(01)012] R. D. van der Mei, J. L. van den Berg, R. Vranken, and B. M. M. Gijssen. Analysis of a Flow Level Model for TCP Behavior in Case Of Priority Queueing. Technical Report 279TD(01)012, COST-279, 2001.
- [TD(01)013] Ilkka Norros and Petteri Mannersalo. A Most Probable Path Approach to Queueing Systems with General Gaussian Input. Technical Report 279TD(01)013, COST-279, 2001.
- [TD(01)014] Philippe Olivier and Nabil Benameur. Flow Level IP Traffic Characterization. Technical Report 279TD(01)014, COST-279, 2001.
- [TD(01)015] R. Schulcz and S. Imre. Handover Support in WATM Networks. Technical Report 279TD(01)015, COST-279, 2001.
- [TD(01)016] Sándor Molnár, Zs. Kenesi, A. Veres, and G. Vattay. Self-Similarity Propagated over the Internet. Technical Report 279TD(01) 016, COST-279, 2001.
- [TD(01)017] W. Burakowski, A. Bak, F. Ricciato, S. Salsano, and H. Tarasiuk. Traffic Handling in AQUILA QoS IP Network. Technical Report 279TD(01) 017, COST-279, 2001.
- [TD(02)001] J. Bachmann, M. Matthes, O. Drobnik, and U. R. Krieger. Traffic Mobility in Wireless IP-Networks Subject to Mobility- and Resource-Management Conditions. Technical Report 279TD(02)001, COST-279, 2002.
- [TD(02)002] H. Reittu and I. Norros. On Large Random Graphs of the “Internet Type”. Technical Report 279TD(02)002, COST-279, 2002.
- [TD(02)003] I. Norros. On a Gaussian Queue with Bandwidth Allocation by Prediction. Technical Report 279TD(02)003, COST-279, 2002.

- [TD(02)004] Sándor Molnár and Gyorgy Terdik. A Monofractal Model for Network Traffic. Technical Report 279TD(02)004, COST-279, 2002.
- [TD(02)005] Jens Milbrandt, Dirk Staehle, Stefan Köhler, and Les Berry. Decomposition of Large IP Networks for Routing Optimization. Technical Report 279TD(02)005, COST-279, 2002.
- [TD(02)006] D. Staehle, K. Leibnitz, K. Heck, P. Tran-Gia, B. Schröder, and A. Weller. Analytical Characterization of the Soft Handover Gain in UMTS. Technical Report 279TD(02)006, COST-279, 2002.
- [TD(02)007] Daniel Zaragoza and Carlos A. Carvalho Belo. Queuing Behavior of Multiplexed General ON-OFF Sources: An Engineering Perspective. Technical Report 279TD(02)007, COST-279, 2002.
- [TD(02)008] Rene Boel and Stijn De Vuyst. Prediction Based Resource Allocation, a Simulation Experiment. Technical Report 279TD(02)008, COST-279, 2002.
- [TD(02)009] Ezhan Karasan and Mustafa Arisoylu. Subnetwork Partitioning and Section Restoration in Translucent Optical Networks. Technical Report 279TD(02)009, COST-279, 2002.
- [TD(02)010] R. Vranken, R. D. van der Mei, R. E. Kooij, and J. L. van den Berg. Performance of TCP with Multiple Priority Classes. Technical Report 279TD(02)010, COST-279, 2002.
- [TD(02)011] Michel Mandjes, Debasis Mitra, and Werner Scheinhardt. Models of Network Access Using Feedback Fluid Queues. Technical Report 279TD(02)011, COST-279, 2002.
- [TD(02)012] Richard J. Boucherie and Aljaz Ule. Adaptive Dynamic Channel Borrowing in Road-Covering Mobile Networks. Technical Report 279TD(02)012, COST-279, 2002.
- [TD(02)013] Markus Fiedler and Gerhard Haßlinger. Waiting time quantiles for the Gaussian voice traffic model. Technical Report 279TD(02)013, COST-279, 2002.
- [TD(02)014] Alberto Medina, Nina Taft, Kavé Salamatian, Supratik Bhattacharyya, and Christophe Diot. Traffic Matrix Estimation, State of Art and New Directions. Technical Report 279TD(02)014, COST-279, 2002.

- [TD(02)015] Kavé Salamatian, Bruno Baynat, and Thomas Bugnazet. Cross Traffic Estimation by Loss Process Analysis. Technical Report 279TD(02)015, COST-279, 2002.
- [TD(02)016] Remco Litjens M. Sc. and Hans van den Berg. Fair Adaptive Scheduling in Integrated Services UMTS Networks. Technical Report 279TD(02)016, COST-279, 2002.
- [TD(02)017] Guy Latouche and Marie-Ange Remiche. A MAP-based Poisson Cluster Model for Web Traffic. Technical Report 279TD(02)017, COST-279, 2002.
- [TD(02)018] Paulo Salvador, Antonio Pacheco Pires, and Rui Valadas. Multiscale Fitting Procedure for Markov Modulated Poisson Processes. Technical Report 279TD(02)018, COST-279, 2002.
- [TD(02)019] F. Delcoigne, A. Proutiere, and G. Regnie. Modelling the Integration of Streaming and Elastic Traffic. Technical Report 279TD(02) 019, COST-279, 2002.
- [TD(02)020] Ilkka Norros and Hannu Reittu. On Large Random Graphs with Infinite Variance Pareto Degree Distribution. Technical Report 279TD(02)020, COST-279, 2002.
- [TD(02)021] Markus Fiedler, Patrik Carlsson, Kurt Tutschku, Stefan Chevul, and Arne A. Nilsson. Obtaining reliable bit rate measurement in SNMP-managed networks. Technical Report 279TD(02)021, COST-279, 2002.
- [TD(02)022] J. Virtamo, E. Kuumola, and J. Resing. Joint Distribution of Instantaneous and Averaged Queue Length in an M/M/1/K System. Technical Report 279TD(02)022, COST-279, 2002.
- [TD(02)023] J. van Leeuwen, S. Aalto, and J. Virtamo. Load Balancing in Cellular Networks Using First Policy Iteration. Technical Report 279TD(02) 023, COST-279, 2002.
- [TD(02)024] Esa Hyttiä and Elena Siren. Delay Line Configurations in Optical Burst Switching with JET Protocol. Technical Report 279TD(02) 024, COST-279, 2002.
- [TD(02)025] Petteri Mannersalo. Some notes on prediction of teletraffic. Technical Report 279TD(02)025, COST-279, 2002.

- [TD(02)026] Nail Akar, Muammer Atik, Inanç Dogru, Ibrahim Hökelek. ATM Multipath Traffic Engineering Using Differentiated ABR. Technical Report 279TD(02) 026, COST-279, 2002.
- [TD(02)027] M. Ajmone Marsan, M. Franceschinis, E. Leonardi, F. Neri, and A. Tarello. Instability Phenomena in Underloaded Packet Networks with QoS Schedulers. Technical Report 279TD(02)027, COST-279, 2002.
- [TD(02)028] Michaela Meo, Marco Ajmone Marsan, and Cecilia Batetta. Resource Management Policies in GPRS Wireless Internet Access Systems. Technical Report 279TD(02)028, COST-279, 2002.
- [TD(02)029] Wojciech Burakowski and H. Tarasiuk. Admission Control for TCP Connections in QoS IP Network. Technical Report 279TD(02)029, COST-279, 2002.
- [TD(02)030] V. Inghelbrecht, B. Steyaert, S. Wittevrongel, and H. Bruneel. Analysis of the Interdeparture Process in Consecutive Stages of a VoIP Network. Technical Report 279TD(02)030, COST-279, 2002.
- [TD(02)031] Michael Menth and Oliver Rose. Performance Tradeoffs for Header Compression in MPLS Networks. Technical Report 279TD(02) 031, COST-279, 2002.
- [TD(02)032] Erich Plasser. On the Non Linearity of the RED Drop Function. Technical Report 279TD(02)032, COST-279, 2002.
- [TD(02)033] Olav Øserbø Telenor Fou. An Approximative Method to Calculate the Distribution of End-to-End Delay in Packet Networks. Technical Report 279TD(02)033, COST-279, 2002.
- [TD(02)034] Hans van den Berg, Robert J. Kooij, Michel Mandjes, Pasi Lassila, and Robert van der Mei. Refinement of PS Models for TCP Performance. Technical Report 279TD(02)034, COST-279, 2002.
- [TD(02)035] Peixia Gao, Sabine Wittevrongel, and Herwig Bruneel. Queueing Analysis of a Discrete-Time Multiserver Buffer System with Geometric Service Times. Technical Report 279TD(02)035, COST-279, 2002.
- [TD(02)036] Joris Walraevens, Sabine Wittevrongel, and Herwig Bruneel. Performance Evaluation of Optical Packet Switches Using Probability Generating Functions: an Initial Analysis. Technical Report 279TD(02) 036, COST-279, 2002.

- [TD(02)037] Hung Tuan Tran and Thomas Ziegler. Engineering Solution of a CAC Mechanism for Voice Traffic over IP Networks. Technical Report 279TD(02)037, COST-279, 2002.
- [TD(02)038] Oznur Ozkasap and Mine Caglar. Traffic Properties of Scalable Multicast Communication. Technical Report 279TD(02)038, COST-279, 2002.
- [TD(02)039] Wojciech Burakowski and Monika Fudala. PFS Scheme for Forcing Better Service in Best Effort IP Network. Technical Report 279TD(02)039, COST-279, 2002.
- [TD(02)040] Jorma Kilpi. A Portrait of a GPRS/GSM Session. Technical Report 279TD(02)040, COST-279, 2002.
- [TD(02)041] Riikka Susitaival, Jorma Virtamo, and Samuli Aalto. Load Balancing by MPLS in Differentiated Services Networks. Technical Report 279TD(02)041, COST-279, 2002.
- [TD(02)042] Klaus Heck, Dirk Staehle, and Kenji Leibnitz. Diversity Effects on the Soft Handover Gain in UMTS Networks. Technical Report 279TD(02)042, COST-279, 2002.
- [TD(02)043] Francisco-Javier Ramón, José Enríquez, Jorge Andrés, and Antonio Molínes. Multipath Routing with Dynamic Variance. Technical Report 279TD(02)043, COST-279, 2002.
- [TD(02)044] Mihael Mohorcic, Gorazd Kandus, and Alex Svigelj. Traffic Flow Model for Routing Study in Packet-Switched Intersatellite Link Networks. Technical Report 279TD(02)044, COST-279, 2002.
- [TD(02)045] Remco Litjens M. Sc. and Richard J. Boucherie. Performance Aanalysis of Downlink Shared Channels in a UMTS Network. Technical Report 279TD(02)045, COST-279, 2002.
- [TD(02)046] Kurt Tutschku and Markus Fiedler. Application of the Stochastic Fluid Flow Model for bottleneck identification and classification. Technical Report 279TD(02)046, COST-279, 2002.
- [TD(03)001] Remco Litjens M. Sc., Hans van den Berg, Richard J. Boucherie, Frank Roijers, and Maria Fleuren. Performance Analysis of Wireless LANs: an Integrated Packet/Flow Level Approach. Technical Report 279TD(03)001, COST-279, 2003.

- [TD(03)002] Kathleen Spaey, Tom Hofkens, and Chris Blondia. Timescales in Models for Bursty Traffic. Technical Report 279TD(03)002, COST-279, 2003.
- [TD(03)003] Riikka Susitaival and Samuli Aalto. Providing Differentiated Services by Load Balancing and Scheduling in MPLS Networks. Technical Report 279TD(03)003, COST-279, 2003.
- [TD(03)004] Chris Blondia, Olga Casals, Llorenç Cerdà, N. Van den Wijngaert, Gert Willems, and Peter De Cleyne. Low Latency Handoff Mechanisms and Their Implementation in an IEEE 802.11 Network. Technical Report 279TD(03)004, COST-279, 2003.
- [TD(03)005] Herrmann de Meer and Kurt Tutschku. A Measurement Study on Signaling in Gnutella Overlay Networks. Technical Report 279TD(03)005, COST-279, 2003.
- [TD(03)006] José Soler Lucas, Villy Baek Iversen, and Martin Nord. Simulation and Performance Analysis of a GMPLS Lambda Scheduler. Technical Report 279TD(03)006, COST-279, 2003.
- [TD(03)007] Thomas Bonald, Philippe Olivier, and James Roberts. Dimensioning High Speed IP Access Networks. Technical Report 279TD(03) 007, COST-279, 2003.
- [TD(03)008] Paulo Salvador, António Nogueira, and Rui Valadas. Joint Characterization of the Packet Arrival and Packet Size Processes of Multifractal Traffic based on Stochastic L-Systems. Technical Report 279TD(03) 008, COST-279, 2003.
- [TD(03)009] Sándor Imre and Mate Szalay. Reliability Modeling of Tree Topology Micro Mobility Networks. Technical Report 279TD(03)009, COST-279, 2003.
- [TD(03)010] Ezhan Karasan, Nail Akar, Oya Karasan, and Mustafa Pinar. Mesh Topology Design in Overlay Virtual Private Networks. Technical Report 279TD(03)010, COST-279, 2003.
- [TD(03)011] Nail Akar and K. Sohraby. Numerically Stable Solution of Large-Scale Finite Fluid Queues. Technical Report 279TD(03)011, COST-279, 2003.
- [TD(03)012] Andrea Baiocchi and Andrea de Vendictis. Investigating TCP Single Source Behavior in Time-Varying Capacity Network Scenarios. Technical Report 279TD(03)012, COST-279, 2003.

- [TD(03)013] Branka Zovko-Cihlar, Winton Afric, and Sonja Grgic. Interference in Direct Spread Sequence Mobile Communication System. Technical Report 279TD(03)013, COST-279, 2003.
- [TD(03)014] Daniel Z. Lenardic, Branka Zovko-Cihlar, and Mislav Grgic. Analysis of Network Buffering Effects on TCP/IP Protocol Behavior. Technical Report 279TD(03)014, COST-279, 2003.
- [TD(03)015] Koenraad Laevens and Herwig Bruneel. Analysis of a Single-Wavelength Optical Buffer. Technical Report 279TD(03)015, COST-279, 2003.
- [TD(03)016] Ilkka Norros and Hannu Reittu. On the Architecture of the Power-Law Random Graph Model of Internet. Technical Report 279TD (03)016, COST-279, 2003.
- [TD(03)017] Dieter Fiems, Stijn de Vuyst, and Herwig Bruneel. Discrete-Time Analysis of the Gated-Exhaustive Vacation Queue. Technical Report 279TD(03)017, COST-279, 2003.
- [TD(03)018] Dirk Staehle and Andreas Mäder. An Analytic Approximation of the Uplink Capacity in a UMTS Network with Heterogeneous Traffic. Technical Report 279TD(03)018, COST-279, 2003.
- [TD(03)019] Stefan Köhler and Andreas Binzenhöfer. MPLS Traffic Engineering in OSPF Networks — A Combined Approach. Technical Report 279TD(03)019, COST-279, 2003.
- [TD(03)020] Jouni Karvo and Samuli Aalto. Using Multicast or a Combination of Unicast and Broadcast for Transmitting Popular Content. Technical Report 279TD(03)020, COST-279, 2003.
- [TD(03)021] Markus Fiedler, Kurt Tutschku, Patrik Carlsson, and Arne A. Nilsson. Identification of performance degradation in IP networks using throughput statistics. Technical Report 279TD(03)021, COST-279, 2003.
- [TD(03)022] R. Litjens, H. van den Berg, and R. J. Boucherie. Throughput Measures for Processor Sharing Models. Technical Report 279TD(03) 022, COST-279, 2003.
- [TD(03)023] Jorma Kilpi. Distributional Properties of GPRS/GSM Session Volumes and Durations. Technical Report 279TD(03)023, COST-279, 2003.
- [TD(03)024] Aleš Švigelj, Mihael Mohorčič, Gorazd Kandus, Andrej Kos, Matevž Pustišek, and Janez Bešter. Consideration of Empirical IP Traffic in

- Traffic Class Dependent Packet-Switched ISL Routing. Technical Report 279TD(03) 024, COST-279, 2003.
- [TD(03)025] Sándor Imre. Dynamically Optimised Chernoff Bound Based CAC for 3G/4G WCDMA Systems. Technical Report 279TD(03)025, COST-279, 2003.
- [TD(03)026] Nabil Benameur, Sara Oueslati, and James W. Roberts. Experimental Implementation of Implicit Admission Control. Technical Report 279TD(03)026, COST-279, 2003.
- [TD(03)027] N. Akar, I. Hokelek, M. Atik, and E. Karasan. A Reordering-Free Multipath Traffic Engineering Architecture for DiffServ-MPLS Networks. Technical Report 279TD(03)027, COST-279, 2003.
- [TD(03)028] Ivan Gojmerac, Thomas Ziegler, Fabio Ricciato, and Peter Reichl. Adaptive Multipath Routing for Efficient Load Balancing in the Internet. Technical Report 279TD(03)028, COST-279, 2003.
- [TD(03)029] Michael Menth, Joachim Charzinski, and Stefan Kopf. Impact of Resilience Requirements on the Performance of Network Admission Control Methods. Technical Report 279TD(03)029, COST-279, 2003.
- [TD(03)030] Guoqiang Hu, Klaus Dolzer, and Christoph Gauger. Does Burst Assembly Really Reduce the Self-Similarity? Technical Report 279TD(03)030, COST-279, 2003.
- [TD(03)031] Sandor Szabo, Sándor Imre, and Alexandrosz Burulitisz. On the Accuracy of Mobility Modelling in Wireless Networks. Technical Report 279TD(03)031, COST-279, 2003.
- [TD(03)032] Michela Meo, Marco Ajmone Marsan, Luca Muscariello, Marco Mellia, and Renato Lo Cigno. A Simple Markovian Approach to Model Internet Traffic at Edge Routers. Technical Report 279TD(03)032, COST-279, 2003.
- [TD(03)033] Detlef Sass, Klaus Dolzer, Stefan Bodamer, and Martin Lorang. Analytic Fluid Flow Approach for Fair Queueing Systems. Technical Report 279TD(03)033, COST-279, 2003.
- [TD(03)034] Hans van den Berg, Michel Mandjes, Remco van de Meent, Aiko Pras, Frank Roijers, and Pieter Venemans. QoS Aware Bandwidth Provisioning in IP Backbone Networks. Technical Report 279TD(03) 034, COST-279, 2003.

- [TD(03)035] Wojciech Burakowski and Andrzej Beben. Premium Message Handling Service in IP QoS Networks. Technical Report 279TD(03) 035, COST-279, 2003.
- [TD(03)036] Ozgur Ozkasap and Mine Caglar. Traffic Characterization of Scalable Multicasting in the Case of a Self-Similar Source. Technical Report 279TD(03)036, COST-279, 2003.
- [TD(03)037] V. Inghelbrecht, B. Steyaert, and H. Bruneel. Burstification Mechanism of an OBS Edge Router: Analytical Analysis of a One- and Two-Threshold Case. Technical Report 279TD(03)037, COST-279, 2003.
- [TD(03)038] Udo Krieger, V. Klimenok, L. Breuer, and A.N. Dudin. The BMAP/G/1 Queue with Feedback Operating in Synchronous Random Environment as a Model for a Telecommunication Channel with Performance Fluctuation. Technical Report 279TD(03)038, COST-279, 2003.
- [TD(03)039] Udo Krieger and N. M. Markovitch. On-Line Estimation of Heavy-Tailed Traffic Characteristics in Web Data Mining. Technical Report 279TD(03)039, COST-279, 2003.
- [TD(03)040] J. Chazinski, C. Hoogendoorn, K. Schrodi, C. Winkler, and M. N. Huber. Towards Carrier-Grade Next Generation Networks. Technical Report 279TD(03)040, COST-279, 2003.
- [TD(03)041] Peixia Gao, Sabine Wittevrongel, and Herwig Bruneel. Queueing Analysis of Multiserver Buffers with Geometric Service Times and Correlated Input Traffic. Technical Report 279TD(03)041, COST-279, 2003.
- [TD(03)042] Patrik Carlsson, Markus Fiedler, and Anders Ekberg. On an Implementation of a Distributed Passive Measurement Infrastructure. Technical Report 279TD(03)042, COST-279, 2003.
- [TD(03)043] Krzysztof Tworus, Stijn de Vuyst, Sabine Wittevrongel, and Herwig Bruneel. Queueing Analysis of the Stop-and-Wait ARQ Protocol in a Wireless Environment. Technical Report 279TD(03)043, COST-279, 2003.
- [TD(03)044] Wojciech Burakowski and Marek Dabrowski. Assessment of Token Bucket Parameters by On-Line Traffic Measurements. Technical Report 279TD(03)044, COST-279, 2003.
- [TD(03)045] Andrzej Beben and Robert Janowski. Statistical Admission Control for WWW Traffic. Technical Report 279TD(03)045, COST-279, 2003.

- [TD(03)046] Michael Menth, Jens Milbrandt, and Andreas Reifert. Backup Capacity Minimization for Simple Protection Switching Mechanisms. Technical Report 279TD(03)046, COST-279, 2003.
- [TD(03)047] Hans van den Berg, Aloysius Irwan Endrayanto, and Richard J. Boucherie. Characterizing CDMA Downlink Feasibility via Effective Interference. Technical Report 279TD(03)047, COST-279, 2003.
- [TD(03)048] Michel Mandjes, Petteri Mannersalo, Ilkka Norros, and Miranda van Uitert. Most Probable Busy Period Paths in Gaussian Queues. Technical Report 279TD(03)048, COST-279, 2003.
- [TD(03)049] Kurt Tutschku and Phuoc Tran-Gia. A Traffic Profile of the eDonkey Filesharing Service. Technical Report 279TD(03)049, COST-279, 2003.
- [TD(03)050] Oznur Ozkasap. Scalability and Robustness of Pull-based Anti-entropy Distribution Model. Technical Report 279TD(03)050, COST-279, 2003.