Helsinki University of Technology Department of Communications and Networking
Teknillinen korkeakoulu Tietoliikenne- ja tietoverkkotekniikan laitos
Espoo 2008 Report 2/2008

TRAFFIC MATRIX ESTIMATION IN THE INTERNET: MEASUREMENT ANALYSIS, ESTIMATION METHODS AND APPLICATIONS

Ilmari Juva

Dissertation for the degree of Doctor of Science in Technology to be presented with due permission for public examination and debate in Auditorium S4 at Helsinki University of Technology (Espoo, Finland) on the 18th of April, 2008, at 12 o'clock noon.

Helsinki University of Technology
Faculty of Electronics, Communications and Automation
Department of Communications and Networking

Teknillinen korkeakoulu
Elektroniikan, tietoliikenteen ja automaation tiedekunta
Tietoliikenne- ja tietoverkkotekniikan laitos

Distribution:
Helsinki University of Technology
Department of Communications and Networking
P.O.Box 3000
FIN-02015 TKK
Tel. +358-9-451 5300
Fax. +358-9-451 2474

ISBN 978-951-22-9298-1
ISBN 978-951-22-9299-8 (pdf)
ISSN 1797-478X
ISSN 1797-4798 (pdf)

Multiprint Oy
Espoo 2008

TEKNILLINEN KORKEAKOULU
TEKNISKA HÖGSKOLAN
HELSINKI UNIVERSITY OF TECHNOLOGY

HELSINKI UNIVERSITY OF TECHNOLOGY P.O.BOX 1000, FIN-02015 TKK http://www.tkk.fi/	ABSTRACT OF DOCTORAL DISSERTATION
Author: Ilmari Juva	

Name of the dissertation
Traffic Matrix Estimation in the Internet: Measurement Analysis, Estimation Methods and Applications

Date of manuscript 10th of December 2007	Date of the dissertation 18th of April 2008
☐ Monograph	☒ Article dissertation (summary + original articles)

Faculty	Faculty of Electronics, Communications and Automation
Department	Department of Communications and Networking
Field of Research	Teletraffic theory
Opponent	Professor Mikael Johansson (KTH Royal Institute of Technology, Sweden)
Supervisor	Professor Jorma Virtamo (TKK Helsinki University of Technology)

Abstract

In a communication network, the traffic has a source, from which that particular traffic flow originates, and a destination, at which it terminates. Each origin-destination (OD) combination constitutes an OD pair. The knowledge of the amount of traffic of each such OD pair in the network is represented by a *traffic matrix*. The traffic matrix is a required input in many network management and traffic engineering tasks, where typically the traffic volumes are assumed to be known. However, in reality, they are seldom readily obtainable, but have to be estimated. The estimators use as input the available information, namely link load measurements and routing information. Solving the OD-pair traffic loads from these is a heavily underconstrained problem. Thus, it is not solvable unless some extra information is brought into the problem.

In the first part of the thesis we analyze measurements from a backbone link of the Finnish University and Research Network (Funet). We consider first the aggregate traffic on the link and then divide the traffic into OD pairs based on the IP addresses of the packets. The traffic traces are analyzed and the traffic is characterized in order to gain insight into the nature of Internet traffic and to study the validity of assumptions necessary in traffic matrix estimation, such as the Gaussian IID model and the functional relation between mean and variance of the traffic volume.

The second part of the thesis studies traffic matrix estimation. We give a brief overview of the proposed methods and note that the majority of them can be classified into two classes based on the extra information that the methods use. These are either the gravity model class or the class that uses the variance through the mean-variance relation.

We derive analytically the Cramér-Rao bounds for the variance of the maximum likelihood estimator. This makes it possible to analyze the performance bounds for the accuracy that can be achieved by the estimator. We propose two novel methods for traffic matrix estimation. The Quick method, based on link covariances, yields an analytical expression for the estimate and is thus computationally light-weight. The accuracy of the method is compared with that of other methods using second moment estimates by simulation under synthetic traffic scenarios. The Combined method incorporates both sources of extra information. This method is shown in many cases to outperform the current estimation methods that rely only on one or other of the sources.

In the third part of the thesis we study robust load balancing. Many traditional load balancing techniques assume the availability of an accurate traffic matrix. However, robust load balancing takes a different approach, and thus does not typically require knowledge of the traffic matrix. We study the robust method but also introduce a new variant of it where the accuracy of the robust method is improved by using an estimated traffic matrix. In this approach we take account the uncertainty in the estimator's accuracy.

Keywords	Internet measurements, traffic characterization, traffic engineering, traffic matrix estimation, robust load balancing.		
ISBN (printed)	978-951-22-9298-1	ISBN (pdf)	978-951-22-9299-8
ISSN	1797-478X	ISSN(pdf)	1797-4798
Language	English	Number of pages	70 pp.
Publisher	Department of Communications and Networking / Helsinki University of Technology		
Print distribution			

☒ The dissertation can be read at http://lib.hut.fi/Diss/

TEKNILLINEN KORKEAKOULU
TEKNISKA HÖGSKOLAN
HELSINKI UNIVERSITY OF TECHNOLOGY

TEKNILLINEN KORKEAKOULU PL 1000, FIN-02015 TKK http://www.tkk.fi/	VÄITÖSKIRJAN TIIVISTELMÄ
Tekijä: Ilmari Juva	
Väitöskirjan nimi Liikennematriisin estimointi Internetissä: Mittausten analysointi, estimointimenetelmät ja sovellukset	
Käsikirjoituksen jättämispäivämäärä 10. joulukuuta 2007	Väitöstilaisuuden ajankohta 18. huhtikuuta 2008
☐ Monografia	☒ Yhdistelmäväitöskirja (yhteenveto + erillisartikkelit)

Tiedekunta	Elektroniikan, tietoliikenteen ja automaation tiedekunta
Laitos	Tietoliikenne- ja tietoverkkotekniikan laitos
Tutkimusala	Teleliikenneteoria
Vastaväittäjä	Professori Mikael Johansson (Kungliga Tekniska högskolan, Ruotsi)
Työn valvoja	Professori Jorma Virtamo (Teknillinen korkeakoulu)

Tiivistelmä

Tietoliikenneverkossa kulkevalla liikenteellä on tietty lähtöpiste, josta kyseinen liikennevirta on lähetetty, ja määränpää, johon se on matkalla. Jokainen lähtöpiste ja määränpää muodostavat niin sanotun OD-parin. *Liikennematriisi* sisältää jokaisen tällaisen OD-parin liikennemäärän. Monissa verkon suunnitteluun ja liikenteenhallintaan liittyvissä tehtävissä tieto liikennematriisista on välttämätön. Monesti se oletetaan tunnetuksi, mutta todellisuudessa liikennematriisi on harvoin suoraan saatavilla, joten se on estimoitava. Estimaattorit käyttävät saatavilla olevaa informaatioita, eli linkkien liikennemääriä ja reititystauluja. Liikennematriisin ratkaiseminen näillä tiedoilla on alimäärätty tehtävä eikä sitä voida ratkaista, ellei tilanteesta tehdä jotain lisäoletuksia.

Väitöskirjan ensimmäisessä osassa analysoidaan mittauksia Suomen korkeakoulujen ja tutkimuksen tietoverkon (Funet) runkolinkiltä. Tutkimme ensin linkin kokonaisliikennettä ja jaamme sen jälkeen liikenteen OD-pareihin pakettien IP-osoitteiden perusteella. Liikennemittaukset analysoidaan ja liikenne karakterisoidaan, jotta saadaan käsitys Internet-liikenteen luonteesta. Samalla tutkitaan liikennematriisin estimointiongelmaan liittyiven oletuksien, kuten gaussisuuden, riippumattomuuden ja keskiarvo-varianssi-relaation, paikkansapitävyyttä.

Väitöskirjan toisessa osassa tutkitaan liikennematriisn estimointia. Käytössä olevista menetelmistä annetaan katsaus. Huomataan, että suurin osa niistä voidaan jakaa kahteen pääluokkaan sen perusteella, millaista lisäinformaatiota ne käyttävät. Toisen luokan muodostavat gravitaatiomalliin pohjautuvat menetelmät ja toisen mentelmät, jotka hyödyntävät liikenteen varianssia keskiarvo-varianssi-relaation avulla.

Väitöskirjassa johdetaan analyyttisesti Cramér-Rao-raja suurimman uskottavuuden menetelmää käyttävän estimaattorin varianssille. Tämän avulla voidaan johtaa rajat estimaattorien tarkkuudelle. Väitöskirjassa esitetään myös kaksi uutta menetelmää liikennematriisin estimointiin. Quick-menetelmä, joka perustuu linkkien kovarianssin hyödyntämiseen, antaa analyyttisen lausekkeen estimaatille ja on siksi laskennallisesti nopea. Tämän menetelmän tarkkuutta verrataan muihin vastaaviin menetelmiin simulaatiotutkimuksissa synteettisellä aineistolla. Combined-menetelmä yhdistää molemmat yleisesti käytetyt lisäinformaation lähteet ja saavuttaa monessa tapauksessa tarkemman estimaatin kuin vain jompaan kumpaan pohjautuvat menetelmät.

Väitöskirjan kolmannessa osassa tutkitaan robustia kuormantasausta. Perinteiset kuormantasausmenetelmät tarvitsevat liikennematriisin syötteeksi, mutta robustit menetelmät eivät tyypillisesti sitä tarvitse. Väitöskirjassa tutkitaan robustia lähestymistapaa. Siihen pohjautuen esitetään uusi menetelmä, jossa robustin lähestymistavan tarkkuutta parannetaan käyttämällä hyväksi liikennematriisin estimaattia, mutta ottaen samalla huomioon estimaatin tarkkuuteen liittyvä epävarmuus.

Avainsanat	Mittaukset Internetissä, liikenteen karakterisointi, liikenteenhallinta, liikennematriisin estimointi, robusti kuormantasaus.		
ISBN (painettu)	978-951-22-9298-1	ISBN (pdf)	978-951-22-9299-8
ISSN	1797-478X	ISSN(pdf)	1797-4798
Kieli	englanti	Sivumäärä	70 s.
Julkaisija	Tietoliikenne- ja tietoverkkotekniikan laitos / Teknillinen korkeakoulu		
Painetun väitöskirjan jakelu			
☒ The dissertation can be read at http://lib.hut.fi/Diss/			

PREFACE

The research that resulted in this thesis started in 2003 when I had the opportunity to start working as a research scientist in the Networking Laboratory after finalizing my Master's Thesis. The work has been done in research projects IRoNet, funded by TEKES, and projects FIT and Fancy, both funded by the Academy of Finland, as well as the Euro-NGI Network of Excellence, which also funded my visit to ENST Bretagne in Brest. In addition, a personal scholarship was received from the Nokia Foundation.

First and foremost, I would like to thank my supervisor Professor Jorma Virtamo. I am grateful for the opportunity to work in the Networking Laboratory and the continued support and all the helpful advice. My instructor Dr. Samuli Aalto and my co-author Dr. Riikka Susitaival also deserve a big thank you for their considerable help and co-operation in my research.

I would also like to thank Professor Sandrine Vaton for all her efforts in our joint research and for hosting my visit to ENST Bretagne, Markus Peuhkuri for his valuable help with traffic measurements, Paola Bermolen for the work she did for our joint paper and Dr. Pirkko Kuusela for instructing me in the beginning of my work.

My thanks to all the employees of Networking Laboratory I have worked with over the years for the great working atmosphere.

Finally, I would like to thank my friends and family, especially my parents.

CONTENTS

LIST OF PUBLICATIONS

[1] Ilmari Juva, Riikka Susitaival, Markus Peuhkuri, and Samuli. Aalto. Traffic characterization for traffic engineering purposes: Analysis of Funet data. In *NGI 2005*, Rome, Italy, 2005.

[2] Riikka Susitaival, Ilmari Juva, Markus Peuhkuri, and Samuli Aalto. Characteristics of origin-destination pair traffic in Funet. *Telecommunication Systems*, 33:67 – 88, 2006.

[3] Ilmari Juva, Riikka Susitaival, Markus Peuhkuri, and Samuli Aalto. Effects of spatial aggregation on the characteristics of origin-destination pair traffic in Funet. In *Proceeding of NEW2AN*, St.Petersburg, Russia, 2007.

[4] Paola Bermolen, Sandrine Vaton, and Ilmari Juva. Search for optimality in traffic matrix estimation: A rational approach by Cramér-rao lower bounds. In *Proceedings of NGI 2006*, Valencia, Spain., Apr. 2006.

[5] Ilmari Juva, S andrine Vaton, and Jorma Virtamo. Quick traffic matrix estimation based on link count covariances. In *ICC 2006*, Istanbul, Turkey, 2006.

[6] Ilmari Juva. Sensitivity of traffic matrix estimation techniques to their underlying assumptions. In *ICC 2007*, Glasgow, UK, 2007.

[7] Ilmari Juva. Robust load balancing. In *Globecom 2007*, Washington D.C., United States, 2007.

[8] Ilmari Juva. Robust load balancing in wireless networks. In *EuroFGI Workshop on IP QoS and Traffic Control*, Lisbon, Portugal, 2007.

1 INTRODUCTION

1.1 Background

Originally the Internet was based on the idea of best effort delivery. This allowed simple network nodes and efficient use of the network. In the current Internet there are more and more applications, such as videostreaming, that require large amounts of bandwidth or quality-of-service guarantees to work properly. This calls for more control over the network so that these requirements can be met. Also, as peer-to-peer applications are used for downloading of large files such as movies or computer programs, these tend to use whatever amount of bandwidth is available, no matter how abundantly bandwidth is provisioned in the network.

Even in a network that is overprovisioned so that bandwidth is, on average, sufficiently large, there will still be congestion in some parts of the network. Some links at some moment may get congested by a change in traffic, due to random fluctuations or sudden interest in a particular web page by a large number of users. Also, link failures and subsequent rerouting of traffic can cause abnormally high traffic volumes in some other parts of the network. Hence, a need to control the traffic in the network in order to use the resources evenly in the network has arisen.

Traffic engineering [AMA+99, KKY03] is used to avoid and deal with congestion in the network. If the whole network is congested, the reason probably is insufficient capacity. If only some parts of the network are congested, the reason most likely is non-optimal allocation of resources in the network. This situation emerges when the dimensioning of the network is no longer appropriate for current traffic volumes. The traffic has perhaps changed due to a change in customer behavior or emergence of new services, or the link capacities have been changed by new infrastructure being built by the operator, or even a link failure somewhere in the network. Traffic engineering techniques can be used to alleviate the congestion.

Congestion in the network can be addressed in several ways. One approach is the use of admission control, where restrictions are used in allowing traffic to enter the network. Quality-of-service guarantees can be given to traffic flows to ensure that they get sufficient bandwidth despite the congestion in the network. This approach concentrates on the service of individual users, instead of on the overall resource usage in the network.

Another approach is load balancing. Specifically in the cases where the whole network is not congested, but only parts of it are, load balancing is a useful method to deal with the congestion. This is a common traffic engineering task where the traffic is routed in an optimal way according to the input information and a performance criterion that describes the level of congestion in the network. Load balancing aims at a situation where some of the traffic contributing to the congestion is routed through another path of the network so that it does not travel through the congested links anymore. Thus, the traffic is distributed over the network in a balanced way with regard to the resources available.

Traffic engineering applications typically require information about traffic volumes on the links and routing information. In addition to this, often the traffic matrix is required as an input for load balancing and dimensioning tasks. The traffic matrix, or demand matrix, gives the traffic volumes between each pair of nodes in the network. However, in current IP networks the traffic matrix is not readily available. Instead it has to be inferred by mathematical estimation techniques. This gives rise to the need for the traffic matrix estimation field of research, where the aim is to develop estimators that provide sufficiently accurate traffic matrices when direct measurement of end-to-end traffic streams is not possible. The measured link loads serve as a starting point for these estimators and are combined with some other relevant information about the traffic to yield an estimate for the traffic matrix.

1.2 Load Balancing

A traditional approach to load balancing would be to rely on the traffic matrix in the optimization problem. As traffic is moved from congested links to other links to alleviate congestion in that part of the network, the traffic matrix is typically needed. The load balancing problem can be formulated as an optimization problem that gets the traffic matrix as an input and yields the optimal routing. The goal is to find optimal value for the objective function. The maximal link utilization in the network is a good choice for objective function. Minimizing the relative link load on the most heavily congested link leads to a situation where there is free capacity on all links, so that an increase of traffic can be accommodated. The decision variable in the optimization problem is the routing matrix. An obvious constraint is the non-negativity of the routing matrix. Another constraint is the requirement that all the traffic gets routed from the origin node to the destination node of the traffic flow in question. Thus the optimization problem can be formulated as follows.

Problem 1 (Load balancing) Given a traffic matrix, find the non-negative routing matrix that minimizes the maximal relative link load and gets all the traffic routed from their origin to their destination.

In wireless mesh networks [AW05] the situation is slightly more complex, which allows also for more flexibility in the optimization. The link capacities are not completely set by the infrastructure, but can be adjusted within certain limits. This is achieved through scheduling. The restricting factor is that nodes that are close to each other interfere with each others' transmission. In the Boolean interference model, when one link is transmitting, all the links within the interference range from this link cannot be active at the same time. Thus, we have several transmission modes, each of which constitutes one possible combination of simultaneously transmitting nodes. By giving more time to some modes at the expense of other's, we can change the capacities of the links, as links belonging to the modes in question can now send more traffic. This adds a new dimension to load balancing, since, in addition to routing, the scheduling of the transmission modes is also a decision variable in the optimization problems. The objective function remains the same, but the scheduling affects the capacities and hence the objective function, since we are optimizing the relative link load.

Problem 2 (Load balancing in wireless network) Given a traffic matrix, find the non-negative routing matrix and transmission schedule vector that minimizes the maximal relative link load and gets all the traffic routed from their origin to their destination.

The traditional approaches described above need the traffic matrix as an input. If the traffic matrix is not available, a different approach can be taken, namely the robust approach. Robust methods are typically used when there is significant uncertainty involved in the problem. If, for instance, the true value of a parameter is unknown, instead of finding a single solution to optimize performance using an estimated value of the parameter, the robust approach finds a solution that works reasonably well over the whole set of possible parameter values. In the present case the traffic matrix is the unknown parameter. A robust load balancing approach can be used taking into account a larger set of possible traffic matrices. The only certainty is the link load measurements, which limit the range of possible traffic matrices to those that are in line with the measurements. The goal is, therefore, to achieve a situation where the optimized routing yields performance that is satisfactory for all the possible traffic matrices that are in line with the measured link loads. The robust problem can be formulated as follows.

Problem 3 (Robust load balancing) Given a set of traffic matrices, find the non-negative routing matrix that minimizes the worst case maximal relative link load over all the traffic matrices, and gets all the traffic routed from their origin to their destination.

1.3 Traffic Matrix

In a communication network, the traffic that transits through the network has a source where that particular traffic flow enters the network and a destination where it exits the network. We can say that the traffic originates from its source and terminates at its destination. These origin and destination points may be links, routers or so called points of presence (POP), depending on the situation and scope within which we look at the network. An origin and a destination point together constitute an origin-destination (OD) pair. The knowledge of the amount of traffic in the network is represented by the traffic matrix, whose elements give the traffic volumes of different OD-pairs. Thus, the traffic matrix describes the traffic of the network.

There is some ambiguity in the definition of the term traffic matrix. If we have only a single measurement, it can be interpreted to be the traffic matrix. On the other hand, in the case where there is a time series of values available, the measurements can be interpreted as samples from a stochastic variable whose expected value is the traffic matrix.

It is widely recognized that traffic matrices accurately representing the traffic demands in the network are crucial for dimensioning and traffic engineering. For the traditional load balancing problem, for example, the traffic matrix is used as an input parameter and is typically assumed to be known.

Obtaining the traffic matrix requires monitoring of the average bitrate over a chosen measurement interval. Direct measurements of the OD-pair bitrates in IP networks can be made by tools such as Netflow [Cis]. These are not, however, typically available throughout the network, and there would also be a significant overhead from the measurements. In Multi Protocol Label Switching (MPLS) [RVC01] networks it is also possible to obtain OD-pair measurement information.

Most commonly only the link load measurements are available. They can be obtained through the Simple Network Management Protocol (SNMP) [CFSD90], which provides the average traffic load on the links, typically with a five-minute measurement interval. Thus, the traffic matrix is not generally easily obtainable in current IP networks. It has to be estimated from the information that is available to us in the network. This easily obtainable information typically consists of routing tables in addition to the link load measurements.

1.4 Traffic Matrix Estimation

The OD-pair traffic loads and link loads are related through the routing matrix. Given the OD-pair traffic loads and the routing matrix, the link loads can be calculated in a straightforward manner. However, as the number of OD pairs in any realistic network is many times larger than the number of links, the reverse problem of inferring the OD-pair traffic from link traffic measurements is heavily underconstrained. This means that estimates are impossible to obtain with just the information that is readily available. Some extra information need to be brought into the problem to make the system identifiable.

While there are several estimation methods proposed in the literature, most of them can be classified under two main approaches based on the two commonly used assumptions that provide the extra information. These are the gravity model assumption [KW95] and the mean-variance relation [Var96, CDWY00]. The third approach is to use direct measurements in the estimation [SLT+05] if they are available.

The gravity model assumes that traffic between two nodes is proportional to the total traffic originating from the source node and terminating at the destination node. This assumption states that no nodes in the network communicate with each other more than their total traffic would allow us to assume. If the assumption holds we can get an accurate prior estimate using just the total traffic leaving from and terminating at each node. This can then be used together with the link counts to yield the final estimate.

In the mean-variance relation the assumption is that the variance of a traffic flow is dependent on the mean volume of that traffic flow through a power-law function. This implies that larger traffic flows also have larger variations in the traffic. Thus, we would be able to use the OD-pair variances in the process of estimating the mean of the OD-pair traffic. If a time series of measurements is available, it is possible to obtain the sample covariances for the link loads. Starting from these, it is then possible to solve the OD-pair variances, which can be used to obtain estimates of the mean of the OD-pair traffic based on the mean-variance relation. Both the first moment (the actual link loads) and the second moment (the covariance of the link loads) of the measurements are utilized. Thus, methods of this kind are called *second moment* methods.

As mentioned above, there is also a third group of estimation methods. These so-called *third generation* methods are, however, fundamentally different from the two classes described above, which start from the assumption that direct measurements are not available. Third generation methods, on the other hand, are based on the possibility of obtaining direct measurements when necessary to calibrate the methods. The problem to be solved then becomes a tradeoff between minimizing the overhead by minimizing the time of measuring while still keeping the estimation error sufficiently small.

Traffic measurements and traffic characterization are tightly linked with the traffic matrix estimation problem. In order to be able to perform traffic engineering tasks such as load balancing, we obviously first need to be able to have measured information from the network telling us the state of the network. But traffic measurements are also an integral part of traffic engineering in another way, not just as input data for the estimators and algorithms. Detailed direct measurements can be used in an offline analysis for characterizing the Internet traffic in general. In particular they can be used for testing the hypotheses related to the traffic matrix estimation problem, such as validity of the gravity model or existence of the mean-variance relation.

1.5 Outline of the Thesis

In this thesis we study the traffic matrix estimation problem. The thesis consists of three parts: traffic measurements and background assumptions, estimation methods and specific traffic engineering applications. Each part is presented in its own chapter. The first part studies measurements obtained from an Internet backbone link. We characterize the traffic in order to gain insight into the nature of Internet traffic and validate some key assumptions involved in traffic matrix estimation. The second part develops estimation methods and statistical performance bounds for estimation accuracy. We also compare the different classes of estimation methods through a simulation study. The last part concentrates on robust load balancing.

The structure of the thesis is as follows. In Chapter 2 we study the measurements from a backbone link of the Finnish University and Research Network (Funet) in different aggregation levels in order to characterize the traffic with traffic engineering purposes in mind. In particular, the measurement data is used to characterize the nature of the OD-pair traffic and specifically to study how certain key assumptions in traffic matrix estimation actually hold in real Internet traffic traces. First, we study the Gaussian IID model, which states that stochastic fluctuations of the traffic can be modeled by the Gaussian distribution and that consecutive measurements are not correlated with each other. We also study cross-correlation between origin-destination pairs. Finally, we explore the existence of a functional relation between the mean and the variance of the OD-pair traffic, as such a relation is absolutely critical to certain estimation techniques.

In Chapter 3 we study the traffic matrix estimation methods. After a review of the existing methods we derive the statistical Cramér-Rao bounds for the well-known maximum likelihood estimation approach. The bound states the lowest variance that can be obtained for an estimator and thus is useful for evaluating performance bounds. We also use this result to propose a method for finding optimal places for direct measurements. A novel quick estimation method is proposed. This uses the same framework as the maximum likelihood estimator (MLE), but makes a tradeoff by sacrificing some of the accuracy in order to be several times quicker than the MLE approach. Finally we divide the existing methods into two main classes based on the assumptions they use in the estimation, and compare the accuracy of these two groups of estimators in situations where their assumptions do not hold exactly.

In Chapter 4 we propose methods for load balancing using a robust approach. In robust load balancing a traffic matrix is typically not used. Link load measurements limit the possible traffic matrices to those that are in line with the measurement results. The load balancing is then performed so that the whole set of possible traffic matrices is taken into consideration. We recognize that only an estimate of the traffic matrix is available, and it has some estimation error. We propose a robust routing scheme that performs well for all plausible traffic matrices, that is, the set of traffic matrices within the confidence interval of the traffic matrix estimate. The Cramér-Rao bounds are used to calculate the standard error and thus the confidence intervals that define a polytope of traffic matrices. This approach to optimize over all plausible, instead of all possible, traffic matrices improves the performance of the robust method. Also, we propose a robust approach for load balancing in wireless networks, where we use a cross-layer approach that optimizes the transmission schedule and routing simultaneously without knowledge of the traffic matrix.

Finally, Chapter 5 summarizes the thesis.

2 CHARACTERIZATION OF TRAFFIC FOR TRAFFIC ENGINEERING PURPOSES

2.1 Introduction

Measurements are necessary for traffic engineering and traffic matrix estimation. The use of the measurements is twofold. First, we need the current link measurements to get an idea of the traffic in the network and data with which to begin the inference or optimization tasks. The link counts are available through the Simple Network Management Protocol [CFSD90]. Second, we need direct measurements for off-line analysis of the traffic and its nature. This is a much broader area for research purposes, and has been given a lot of attention. We concentrate on traffic characterization for traffic engineering purposes, and especially study the assumptions used in traffic matrix estimation.

Various assumptions are commonly made in order to solve the traffic matrix estimation problem. In this section we study the traffic from the Finnish University and Research Network (Funet). The goal is to discover the basic characteristics of Internet traffic and to test the validity of assumptions made in traffic matrix estimation techniques.

For statistical approaches, it is integral to assume a distribution that the unknown traffic volumes are following. The classic method for telephone networks was to use the Poisson distribution. However, there is lot of evidence that this is not applicable to data networks. In particular for short time scales more complex models are needed. For aggregates over longer time scales, however, a simple Gaussian model might be a sufficient approximation. Another common assumption is the independence of consecutive measurements. Together these constitute the Gaussian IID model.

The mean-variance relation assumes a functional relationship between the mean and the variance of an origin-destination (OD) pair. We differentiate between two separate terms that are not to be confused with one another. First, by *spatial relation* we mean a situation where we consider the relation between OD pairs or links. That is, it is studied whether the variance of an OD pair is larger for the OD pairs that have larger traffic volumes. This is a key assumption, as it is the basis of a family of traffic matrix estimation methods. By *temporal relation* we mean that the variance of a particular OD pair's traffic at a given time is related to the volume of the traffic at that time. That is, when there is more traffic, the variation is higher.

2.2 Review of related work

IP Traffic characteristics

There is a vast literature considering traffic characteristics of modern data networks. Leland et al. [LTWW94, LW91] studied Ethernet traffic characteristics of the well-known Bellcore measurements and found the traffic to be bursty and self-similar in nature, meaning that variability is seen on a wide range of timescales. Aggregating traffic does not smooth the trace, instead the aggregate also exhibits similar bursty behavior as the smaller flows.

Paxson and Floyd [PF95, Pax97] discovered self-similarity also in wide area networks, and conclude that Poisson-model while valid for user session arrivals, cannot be used for modeling packet arrival.

Park et al. [PKC96, Par97] find a causal relationship between the self-similarity of traffic and heavy-tail distribution of file size. They also note that in TCP traffic this long-range dependency is visible, but in non-flow-controlled traffic (UDP) little of the self-similar nature induced by the heavy tail distribution is preserved in the aggregate.

Crovella et al. [CB96, CTB98] show that heavy-tailed file size distribution is also prevalent in WWW based traffic, and responsible for the self-similar traffic behavior.

Models to capture the self-similar nature of the traffic in communication networks has been proposed since, starting from the Brownian motion model by Norros [Nor94, Nor95].

Cao et al. [CCLS01] as well Molnar and Dang [MD00] point out that the traffic is non-stationary, and that this has to be taken into account along with the self-similar nature and long-range dependency in order to obtain reliable estimations and conclusions of the characteristics of the traffic.

Karagiannis et al. [KMFB04] question the validity of the decade old classic results, as both capacity and number of hosts in internet has largely increased since then. They propose a non-stationary time-dependent Poisson model that models their newer data set well, and exhibits long range dependence on longer time scales.

On longer time scales, Central Limit Theorem would let us assume that the aggregated traffic could be approximately Gaussian. Kilpi and Norros [KN02] study the aggregation of traffic in both space (number of flows) and time. They use a simple correlation statistic based on the NQ-plot to identify aggregation levels where the traffic can be considered Gaussian. They note that both types of aggregations are needed to yield approximately Gaussian traffic.

Van de Meent et al. [vdMMP06] also studied the approximate Gaussian behavior of traffic. They discover that the number of users required for traffic aggregates to be Gaussian, and found that “a few tens of users” is typically sufficient. They also studied the NQ-plot correlation method, and found that although simple, it is sufficiently accurate to determine the validity of the approximate Gaussian assumption.

Origin-Destination pair traffic characteristics

Fewer papers have studied the origin-destination based traffic characteristics, perhaps due to the relative unavailability of appropriate data sets. This kind of analysis is, however, the basis for traffic matrix estimation, and thus many traffic engineering applications.

Feldman et al. [FGL+01] characterize point-to-multipoint traffic and find that a few demands account for 80% of total traffic and the traffic volumes follow Zipf’s law. They also point out that daily profiles of the greatest demands also vary significantly from each other.

Bhattacharyya et al. [BDJT01] characterize Point of Presence-level (POP) and access-link level traffic dynamics. They find that there are huge differences in the traffic volumes of the demands. In addition, the larger the traffic volume of an egress node, the larger also the variability of the traffic during the day.

Lakhina et al. [LPC+04] present an interesting structural analysis of the traffic demand matrix, based on sampled flow data from backbone networks in Europe and the US. By applying Principal Component Analysis (PCA), they demonstrate that each OD flow can be well approximated by a linear combination of a small number of so called eigenflows. In addition they observe that these eigenflows fall into three categories: *deterministic* exhibiting strong diurnal periodicity, *spike eigenflows* with clear outliers, and *noise eigenflows* with a nearly Gaussian marginal distribution.

Soule et al. [SNC+04] observe from sampled flow data, collected from a commercial Tier-1 backbone, that large and medium size OD flows contain (at least) two sources of variability. This coincides with the eigenflows found by Lakhina et al., as the OD pairs have a deterministic cyclo-stationary diurnal patterns along with a noise term with zero mean.

Cao et al. [CDWY00] propose a moving IID Gaussian model, consisting of a deterministic term capturing the possible cyclo-stationary diurnal pattern and a randomly fluctuating term.

Mean-variance relation

Vardi [Var96] proposed a Poisson model for traffic matrix estimation. This would, in particular, imply that the variance of the traffic equals the mean of the traffic. Thus, in essence, Vardi was the first to suggest a dependence between the size and variability of the OD pair traffic. The Poisson model, however, was too restrictive.

Cao et al. [CDWY00] propose a power-law functional mean-variance relation, where the variance is proportional to the power of the mean.

$$D^2[X_n] = \phi E[X_n]^c. \tag{2.1}$$

They concluded that a quadratic power law with $c = 2$ is a reasonable fit for the local area network in their study to justify its use.

Morris and Lin [ML00] find a linear relationship between the variance and the mean ($c = 1$) for Web traffic. They base the statement on traffic traces from Harvard's campus network (100 Mbps Ethernet) and a local network at Lucent (Ethernet). The time scale used in this study is $\Delta = 0.1$ s.

In a recent study, Gunnar et al. [GJT04] confirm the validity of the mean-variance relationship based on data traces from a global operator's backbone, and give values $c = 1.5$ and 1.6.

Soule et al. [SNC+04] also discover a value of similar size ($c = 1.56$) but point out that the fit of the relation is poor in their case.

Concerning the temporal relation, results in literature are more in line with each other.

Medina et al. [MTS+02] reported that while the power-law mean-variance relationship (2.1) seems to hold, the exponent c varies remarkably from one link to another within bounds $c \in [0.5, 4.0]$. These observations are related to data collected from a tier-1 backbone with time scale $\Delta = 1$ s.

Soule et al. [SNC+04] note that the parameter values as well as the validity of the relation vary dramatically from OD pair to OD pair, with the parameter being in range $c \in [1.0, 4.0]$.

2.3 Contribution of the Thesis

Most of the measurement studies in literature concentrate on small time scales, from microseconds to milliseconds and seconds. The longer time scales, from seconds to minutes and hours, have not been explored that much. However, traffic engineering and traffic matrix estimation typically use the SNMP link load measurements, which is a five minute aggregate of the link traffic. Therefore we are interested in traffic characteristics on the time scale of minutes, instead of sub-second time scales.

In Publications 1, 2 and 3 we study measurements from Finnish University and Research Network (Funet) to this effect. The work is based on measurements made in the Funet network in 2004 and 2006.

Publication 1 introduces the Funet measurements. We study the characteristics of one link's aggregate traffic. The Gaussian assumption as well the IID (independent and identically distributed) assumption are studied at different aggregation levels. The temporal mean-variance relation is also studied.

In Publication 2 the traffic traces are divided into OD pairs. Thus we are able to study the different OD pairs separately. We identify different types of OD pairs and note that the traffic profiles are very different from each other and from the aggregate. Again, the Gaussian IID model is studied, this time for the OD pair traffic. From this data set we are also able to study the important spatial mean-variance relation.

Publication 3 extends the work in Publication 2 to include different levels of spatial aggregation by using different resolutions to obtain the OD pairs from the aggregate data. We study the patterns of diurnal variation of OD-pair traffic. Due to the different levels of aggregation available we are able to examine how spatial aggregation affects the validity of the Gaussian and mean-variance assumptions.

2.4 Funet Measurements

Funet traces were captured between csc0-rtr and helsinki0-rtr[1] from a 2.5 Gbps STM-16 link using Endance DAG 4.23 cards. The IP addresses on captured packet headers were anonymized preserving prefix, and the headers were stored to disk using flow-based compression [Peu01]. Captured traffic was transferred once an hour to an analysis machine where statistics were calculated. Part of the traces were archived for later analysis, but not all because of large volume of the data (about 10 Mbps average). For this study, bytes transferred each second were calculated.

[1]For details about Finnish university and research network (Funet), see http://www.csc.fi/suomi/funet/verkko.html.en

Of the measured traffic, TCP accounts more than 98 % of bytes transferred. During daytime 10-20 % of TCP traffic is HTTP. There exists also considerable amount of peer-to-peer traffic. Part of data points were missing because of transient errors in data analysis.

We divide the traffic of the link into origin-destination pairs by identifying the origin and destination networks of packets by the left-most bits in the IP address. Let l denote the number of bits in this network prefix, also called network mask. Different levels of aggregation are obtained by changing the prefix length l. The maximum length of the network prefix is 24 bits. With this resolution, there are 2^{24}, or over sixteen million, possible origin networks. On the other hand, with the prefix length $l = 1$ there are only two networks and thus four possible OD pairs.

Our procedure for selecting OD-pairs for further analysis from the original link traffic is the following. Combining both directions, the N most active networks in terms of traffic sent are selected and an $N \times N$ traffic matrix is formed, where $N \leq 100$. This is enough to include all the significant OD pairs. From the obtained traffic matrix at most M greatest OD pairs in terms of sent traffic are selected for further analysis. We select $M = 100$, except in section 2.6, where we use $M = 1000$. Note that for very coarse level of aggregation the number of all OD pairs remains under 100.

The measurements capture the traffic of two days: November 30th 2004 and June 31st 2006, with the main focus being on the first day.

2.5 Gaussian IID Model

We test whether a Gaussian assumption is valid for the Funet traffic. The traffic trace is separated into different components from the link bit counts x_n^Δ,

$$x_n^\Delta = m_n^\Delta + s_n^\Delta z_n^\Delta,$$

where m_n^Δ refers to the moving sample-average, s_n^Δ to the moving sample-standard-deviation, and z_n^Δ to the sample-standardized residual. The averaging was done using a moving window over a one hour period. We concentrate on the stochastic component of the traffic, the standardized residual z_n^Δ.

A good way to evaluate the appropriateness of the Gaussian assumption is the normal quantile (N-Q) plot. The original sample vector x is ordered from the smallest to the largest and plotted against vector a, which is defined as

$$a_i = \Phi^{-1}\left(\frac{i}{n+1}\right) \qquad i = 1, \ldots, n,$$

where Φ is the cumulative distribution function of the normal distribution. The vector a thus contains the normal quantiles, having values from approximately -3 to 3. If the considered data follows the normal distribution, the plot should be linear. Goodness of fit with respect to this can be calculated by the linear correlation coefficient r, and the value r^2 is used as a measure of the fit.

$$r(x, a) = \frac{\sum_{i=1}^{n}(x_i - \overline{x})(a_i - \overline{a})}{\sqrt{\sum_{i=1}^{n}(x_i - \overline{x})^2 (a_i - \overline{a})^2}}.$$

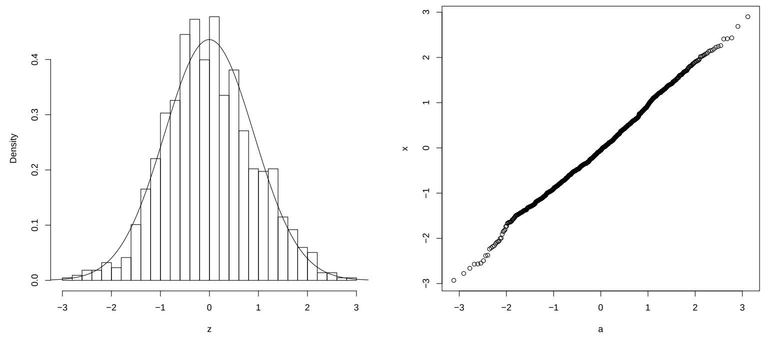

Figure 2.1: Histogram and N-Q test comparing the Funet-data against the normal distribution.

In Figure 2.1 the histogram and Normal-quantile (N-Q) plot of the 5-minute time scale link aggregate are shown comparing them against density function of the normal distribution. The curve does not follow the histogram exactly, but there is a reasonably good fit. The N-Q plot confirms the observed Gaussianity as the goodness-of-fit value is 0.996. For the one second time scale of the same trace it is 0.999.

However, when we study the origin-destination flows in Publications 2 and 3, only a small portion of them seem anywhere close to Gaussian, typically only the larger flows. While it is impossible to set a strict threshold, it seems that in our data majority of the OD pairs with at least 10 Mbps of traffic are fairly Gaussian.

The Gaussianity of each OD pair is evaluated by the N-Q plot goodness of fit, and the value r^2 is used as a measure of the Gaussianity. In Figure 2.2 the size of the OD pair traffic volume (bits per second) is plotted against the goodness of fit value r^2 of the Gaussian assumption. We can see from the figure that the larger flows are always close to Gaussian, with r^2 values easily over 0.90. The largest OD pair with $r^2 < 0.90$ has traffic volume of 17.5 Mbps. The vertical line in the figure is located at 10 Mbps, which seems to be an approximate threshold beyond which an overwhelming majority of the OD pairs have $r^2 > 0.90$. Indeed for many of the OD pairs $r^2 > 0.98$. For OD pairs of size from 1 Mbps to 10 Mbps there is still a lot of Gaussian traffic, while for OD pairs smaller than 1 Mbps no Gaussian behavior is observable.

For the measurement sample to be considered IID, there should not be any significant autocorrelation observable in the stochastic component z_n^Δ. The autocorrelation function is defined as:

$$r_l(k) = \frac{\sum_{i=1}^{T/\Delta - l}(z_{i,k}^\Delta - \bar{z}_k^\Delta)(z_{i+l,k}^\Delta - \bar{z}_k^\Delta)}{\sum_{i=1}^{T/\Delta}(z_{i,k}^\Delta - \bar{z}_k^\Delta)^2},$$

where T/Δ is the size of time series and l is the *lag* .

In Figure 2.3 the autocorrelation of the five minute link aggregate is shown. Clearly there are positive autocorrelations, meaning dependency

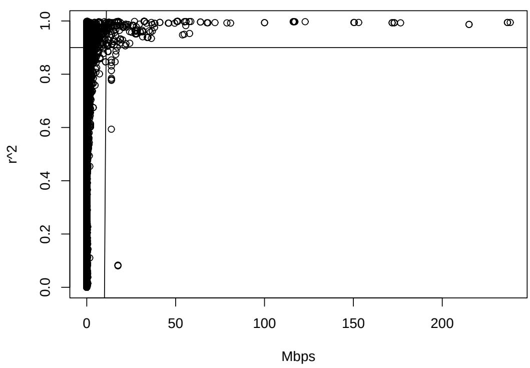

Figure 2.2: Testing Gaussianity: Goodness of fit values r^2 as a function of OD pair traffic volume.

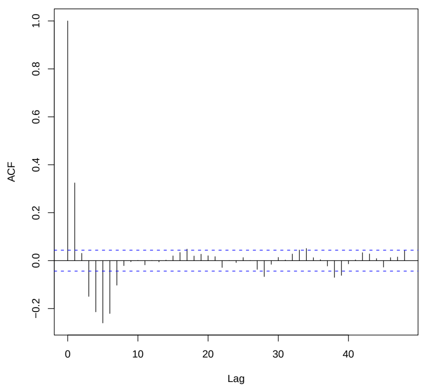

Figure 2.3: Autocorrelation of Funet data with five minute time scale

between consecutive measurements. We notice significant positive value for a lag of five minutes, and after that a set of negative autocorrelation values is clearly observable. It is not until a lag of more than thirty minutes that there is no significant autocorrelation.

In Publication 2 the autocorrelation of OD pairs was studied. We notice that not all of them exhibit long range autocorrelation, but some do.

We also examine whether the OD pairs are independent from each other by studying the dependency between the residual components of the OD pairs. To evaluate this, we have calculated cross-correlation between the residuals $z^{\Delta}_{n,k}$ and $z^{\Delta}_{n,k'}$ of different OD pairs k and k' for the 20 largest

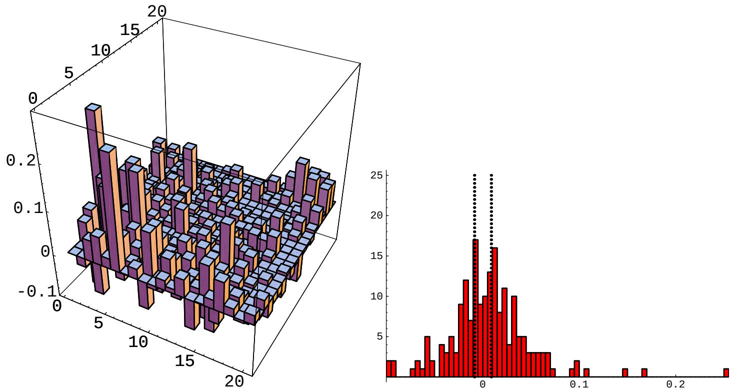

Figure 2.4: Left: Correlation between 20 greatest OD pairs. Right: distribution of correlation coefficients, with 95% confidence interval depicted by dotted lines.

OD pairs:

$$r(k,k') = \frac{\sum_{i=1}^{n}(z_{i,k}^{\Delta} - \bar{z}_{k}^{\Delta})(z_{i,k'}^{\Delta} - \bar{z}_{k'}^{\Delta})}{\sqrt{\sum_{i=1}^{n}(z_{i,k}^{\Delta} - \bar{z}_{k}^{\Delta})^2(z_{i,k'}^{\Delta} - \bar{z}_{k'}^{\Delta})^2}}.$$

The correlation values are presented graphically in Figure 2.4, where we have not considered the terms between the OD pair and itself, which would obviously equal 1.0. The distribution of the various correlation terms in that matrix is also shown in the figure, where the horizontal lines in the figure depict the 95% confidence interval of the hypothesis that correlation would be zero. Clearly there is a large number of statistically significant non-zero values in our data.

2.6 Mean-Variance relation

The commonly used power-law relation between the mean and variance is

$$\mathbf{\Sigma} = \phi \cdot \mathrm{diag}\{\boldsymbol{\lambda}^c\}, \tag{2.2}$$

where $\mathbf{\Sigma}$ is a diagonal matrix, because we assume independence between OD pairs. The relation for an individual OD pair i is

$$\sigma_i^2 = \phi \cdot \lambda_i^c,$$

or for the logarithms

$$\log \sigma_i^2 = c \log \lambda_i + \log \phi.$$

Thus, if the relation held, the points would fall on a line with slope c and intercept $\log \phi$ in the log-log scale. This is a simple linear regression model and we can measure the validity of the mean-variance relation with the linear correlation goodness of fit value r^2 used in the previous section.

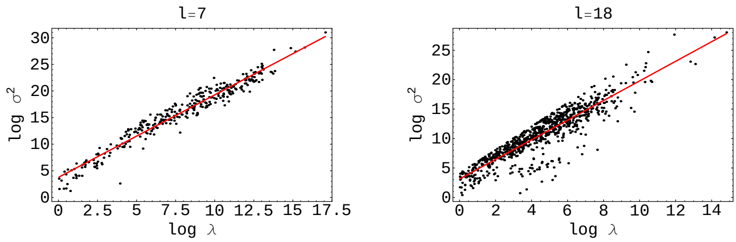

Figure 2.5: Mean variance relation in log-log scale. Left: $r^2 = 0.95$, right: $r^2 = 0.80$.

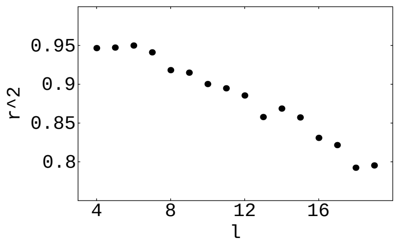

Figure 2.6: Testing mean-variance relation: Goodness of fit values r^2 as a function of prefix length l.

For each prefix length, the mean and the variance are calculated for each one hour period in the 24 hour trace. In Figure 2.5 the values are depicted for one selected hour and two selected prefix lengths, with one point in the plot representing the mean and the variance of one OD pair for that hour. For a longer prefix ($l = 18$) $r^2 = 0.80$, which is in line with previous results. It can be seen that the values deviate significantly more from the regression line making the fit worse. However, for a shorter prefix ($l = 7$), depicted in the same Figure, the fit is much better, about $r^2 = 0.95$.

In Figure 2.6 the average goodness of fits values are shown as a function of the network prefix length l. As the prefix gets longer, there are more OD pairs, with the average size of an OD pair obviously getting smaller. For the longer prefixes the fit of the mean-variance relation is around 0.75 to 0.80. As the resolution gets coarser, the goodness of fit values improve to over 0.90, in some cases as high as 0.95. The OD-pair traffic volumes at these aggregation levels are still less than 100 Mbps, and as the growth is approximately linear as a function of the aggregation level, we may conclude that for larger traffic flows the fit is probably at least as good.

Table 2.1 shows the values of the exponent parameter c with different aggregation levels. It can be said that the parameter stays relatively constant and that the values fall between the results reported for the parameter values in the literature.

We can conclude that there is a clear dependency between the mean-

Table 2.1: Estimates for the mean-variance relations exponent parameter c for different prefix lengths l.

l	4	6	8	10	12	14	16	18	20
c	1.64	1.60	1.66	1.72	1.77	1.75	1.73	1.67	1.71

variance relation fit and the aggregation level. Most importantly, there is a strong functional mean-variance relation for the cases where aggregation level is high.

2.7 Summary and Conclusions

In this chapter we studied the validity of two common traffic engineering assumptions concerning the OD counts: the Gaussian IID model of observed measurements, and the functional relation between mean and variance. Unfortunately we did not have suitable data from the Funet network to enable the gravity model assumption to undergo similar scrutiny. This key assumption cannot therefore be validated nor invalidated by these measurements.

For the link aggregate traffic the normal distribution was found to be a very satisfactory fit, while the small OD pairs are not close to it. We found that the Gaussian assumption holds better when the aggregation level is higher. An approximate threshold, after which all OD pairs are at least fairly Gaussian, would appear to be around traffic volumes of 10 to 20 Mbps. This means that for many traffic engineering and traffic modeling tasks where we consider much larger traffic flows the Gaussian assumption is justified, but it probably cannot be used for cases with smaller traffic volumes due to low aggregation level. Typically traffic matrix estimation is performed on backbone networks, where traffic volumes are larger than the threshold, and the Gaussian distribution therefore seems appropriate.

The autocorrelation function of the Funet traffic shows significant values of autocorrelation for lags as long as half an hour. The strict IID assumption is thus not valid, even for five-minute aggregated measurements, let alone for shorter intervals. Also, the OD pairs are not independent of each other. We found cross-correlation values of up to 0.30, and a large number of smaller, yet still statistically significant non-zero values. Further study is needed to assess whether the autocorrelation and cross-correlation values found are large enough to affect the results when computations are performed assuming independence.

Finally, we validated the spatial power law relation between mean and variance of the OD pairs. The relation seems to hold rather well, particularly with large aggregation levels, where the goodness of fit value is around 0.95. The validation of the mean-variance relation is an essential result concerning many traffic matrix estimation techniques that rely on this very assumption. Our results also show that the exponent parameter remained about constant regardless of the aggregation, and was within the range of values obtained for it in other studies in the literature.

3 TRAFFIC MATRIX ESTIMATION

3.1 Introduction

The traffic matrix gives the volume of traffic traversing a network. Each element of the matrix corresponds to an OD pair in the network. The traffic matrix is a required input for the operator in many network management tasks. Such tasks include for instance routing, traffic engineering problems such as balancing the traffic load in the network as evenly as possible for all links, as well as network capacity dimensioning.

In many proposed traffic engineering methods, knowledge of the underlying traffic volumes is assumed to be known. However, in reality, they are seldom readily available in networks. Instead the traffic matrix has to be typically inferred from the link count measurements. There are several estimation methods proposed to achieve this.

In this thesis we study the problem from several different viewpoints. We derive bounds for the variance for the maximum likelihood methods in order to get an analytical expression for the performance bounds of the method's accuracy. A simulation study is performed to compare the two main classes of estimators. We study the sensitivity of these estimators to their assumptions. Namely, how their accuracy decreases when the assumptions do not hold exactly. We also propose two novel estimators. One is a method that makes a trade-off of slightly more inaccurate estimates for significantly reduced computation time. The other is the first method to combine the two classes of additional information in the estimation process.

3.2 Traffic matrix estimation problem

Each traffic flow in a network originates from some origin, and terminates at some destination. These may be links, routers or so called points of presence (POP), depending on the situation, but in the sequel we will refer to these as nodes. Each origin (or source) node s and destination node d constitute an OD pair. The traffic between the origin and destination of an OD pair is denoted by x_{sd}, which is the element (s, d) of the traffic matrix $\boldsymbol{x}$. For the computational purposes, the traffic matrix is always written as an n-vector $\boldsymbol{x}$, where n is the number of non-zero OD pairs. We refer to the traffic of the ith OD pair by x_i. The vector contains all nonzero elements of the matrix, as zero elements are left out. The unknown traffic matrix is a stochastic variable $\boldsymbol{X}$ and let $\boldsymbol{x}$ be some value of this variable, which are the actual traffic volumes in the network. The traffic matrix, denoted by $\boldsymbol{\lambda}$, is the expected value of $\boldsymbol{x}$.

$$\mathrm{E}[\boldsymbol{X}] = \boldsymbol{\lambda}.$$

The vector $\boldsymbol{\lambda}$ is what we are trying to estimate in traffic matrix estimation, although in many cases also $\boldsymbol{x}$ is estimated.

Available Data

Although direct measurement of traffic matrices is possible with tools like Netflow [Cis], they are typically not available over the whole network, and network wide use of Netflow would be quite expensive. Hence, the information about the OD-pair volumes $\boldsymbol{x}$ is typically not readily available, but has to be estimated. What we do have available are the measurements of the traffic in each link, and also the routing matrix specifying the path each OD pair uses in the network between nodes s and d.

The *link counts*, or link loads, give the measured traffic volumes in each link at a given time. They are denoted by the m-vector $\boldsymbol{y}$. The element y_j of this vector gives the link count on a specific link j. When consecutive measurements are used, we denote the tth set of measurements by vector $\boldsymbol{y}_t$. The link counts are obtained from the measurement data available by the Simple Network Management Protocol (SNMP) [CFSD90]. The attractive feature of SNMP is that it is usually available everywhere in an IP network. However, it has many limitations, such as possible inaccuracy and unreliability as data may be lost in transport. See [CFT$^+$02] for discussion of problems in using SNMP for traffic measurements. Despite the problems, SNMP is the only widespread tool to obtain link count data. The SNMP poller requests periodically each router for the amount of traffic received and transmitted by its interfaces. Measurement periods vary from one minute to few minutes, with five minutes being the typical value.

The *Routing matrix* $\boldsymbol{A}$ is of dimensions $m \times n$ and is usually assumed to be known and fixed in traffic matrix estimation problems. Element $A_{j,i}$ of the routing matrix is 1 if OD pair i uses link j, and 0 otherwise. The routing matrix is obtained from BGP configurations and through OSPF and IS-IS link weight information gathered from the routers.

The Link Count Relation

Should we know the traffic between OD pairs and the routing matrix, the link counts could easily be calculated through the link count equation

$$\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x}. \tag{3.1}$$

In order to be consistent with the link counts the traffic matrix estimate has to satisfy this equation. The link count equation 3.1 will hold for any so-called snapshot of the network. In any particular moment of time the link loads are deterministically obtained from the OD loads by this equation. If we have several measurements available we can use the sample average of $\boldsymbol{y}$ as the expected value of the link loads and obtain an equation including the traffic matrix $\boldsymbol{\lambda}$.

$$\mathrm{E}[\boldsymbol{y}] = \boldsymbol{A}\mathrm{E}[\boldsymbol{x}] \tag{3.2}$$

or

$$\overline{\boldsymbol{y}} = \boldsymbol{A}\boldsymbol{\lambda}. \tag{3.3}$$

The Problem Formulation

The problem setting of traffic matrix estimation can be divided into two different scenarios, depending on the amount of measurements available.

Therefore, there is some ambiguity in the literature whether the term *traffic matrix* refers to OD counts $\boldsymbol{x}$ or their expected value $\boldsymbol{\lambda}$.

If only a single measurement snapshot of the link counts is available, the goal of the estimation is to infer the OD counts $\boldsymbol{x}$ from the link counts $\boldsymbol{y}$. On the other hand, if there is a time-series of several link count measurements $\boldsymbol{y}_t$ $(t = 1, \ldots, T)$ available, the problem usually is to infer the expected value $\boldsymbol{\lambda}$ of the OD counts, although in some cases the goal is to infer also the link count time series $\boldsymbol{x}_t$ $(t = 1, \ldots, T)$.

Problem 1 (Snapshot problem) Given a single set of link counts $\boldsymbol{y}$ and routing matrix $\boldsymbol{A}$, find the traffic matrix $\boldsymbol{x}$ such that conditions $\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x}$ are satisfied.

Problem 2 (Time-series problem) Given independently and identically distributed and locally stationary link counts $\boldsymbol{y}_t$ $(t = 1, \ldots, T)$ and routing matrix $\boldsymbol{A}$, find the traffic matrix $\boldsymbol{\lambda}$ such that conditions $\overline{\boldsymbol{y}}_t = \boldsymbol{A}\boldsymbol{\lambda}$ are realized as close as possible.

Since in any realistic network there are many more OD pairs than links, the problem of solving $\boldsymbol{x}$ or $\boldsymbol{\lambda}$ from $\boldsymbol{A}$ and $\boldsymbol{y}$ is strongly underdetermined and thus ill-posed. This means that accurate explicit solutions cannot be found, as there are infinite number of solutions that satisfy equation (3.1).

To overcome this ill-posed nature of the problem, some type of additional information has to be brought in before the problem can be solved. This might be assumptions about the traffic distribution, additional measurements or some prior knowledge about the traffic matrix. The estimation methods can be classified into different groups based on the nature of the additional information used. In the next section we review the methods proposed in literature.

3.3 Classification of Methods

There are several possible ways to classify the different methods proposed for the traffic matrix estimation problem. In this section we review the different ways of classification.

The key element in making an estimate in an ill-posed situation is the source of the extra information (or side information) brought in to make the problem identifiable. The classification can thus be made based on the extra information used in the method. While in recent years several methods for the traffic matrix estimation problem have been proposed, it is interesting to notice how strikingly few different approaches there really are. The two most common approaches are the gravity model based methods and the second moment methods utilizing the mean-variance relation. While these sources of extra information are not mutually exclusive, it is logical classification in the sense that majority of first and second generation methods are based on the use of one or the other.

A broader classification can be made by the general nature of the extra information. This can be IID model (includes the second moment methods), spatial (includes the gravity methods), temporal or spatio-temporal. Thus the classes here are dependent on the nature of the side information used.

In [SLT+05] the methods are classified into three generations, where the classification is based on the source of the side information. The use of the term generation is inspired by the chronological order in which the different approaches were proposed in literature, but the classification is, in fact, based on the extra information used. First generation methods deploy the second moment statistics of the link counts to obtain the extra information. Second generation methods make use of additional side information other than the SNMP link counts to make the problem identifiable. Most commonly the gravity model is used. Third generation methods assume that direct measurements are available on demand to calibrate the methods. The problem then becomes slightly different in the sense that the traffic matrix is not unknown and the challenge is to build a model that gives accurate real time approximations of the traffic matrix with minimal calibration measurements.

An important distinction between different estimation methods is made also based on what inference strategy they use. The most common strategies are the maximum likelihood approach and the projection approach and its variants, while some methods have tried to use linear programming.

3.4 Review of Proposed Methods

In this section we review the methods proposed in literature. Not all methods fall in a specified class within the classification used here. Therefore we also include a section on other methods.

First Generation IID Model Methods

The first generation IID model methods use the functional mean variance relation to make use of the second moment statistics, namely the covariance matrix, of the link count measurements. To obtain a sample covariance matrix we need a time series of measurements, instead of just a snapshot of the link counts which is sufficient with the gravity based methods.

The work by Vardi in [Var96] is one of the first papers on traffic matrix estimation in computer networks. Vardi was the first to propose a method using the second moments to serve as the additional information to make the system identifiable, and coined the term network tomography, because of the similarities in the problem to medical tomography. A Poisson distribution is assumed, meaning that variance is equal to the mean. This allows the sample covariance matrix of the link counts to be used to estimate the traffic matrix. Thus, basically, the method uses the mean-variance relation with parameters $\phi = 1, c = 1$.

Cao et al. [CDWY00] formulate a maximum likelihood equation using Gaussian distribution and the mean variance relationship. The log-likelihood function conditioned on τ measurements of the linkcounts $\boldsymbol{y}$, can be written as

$$\begin{aligned} l(\boldsymbol{\lambda}, \phi, c|\boldsymbol{y}_1, \ldots, \boldsymbol{y}_\tau) = &-\frac{\tau}{2}\log|\phi\boldsymbol{A}\boldsymbol{\Sigma}'\boldsymbol{A}^{\mathrm{T}}| \\ &-\frac{1}{2}\sum_{t=1}^{\tau}(\boldsymbol{y}_t - \boldsymbol{A}\boldsymbol{\lambda})^{\mathrm{T}}(\phi\boldsymbol{A}\boldsymbol{\Sigma}'\boldsymbol{A}^{\mathrm{T}})^{-1}(\boldsymbol{y}_t - \boldsymbol{A}\boldsymbol{\lambda}), \end{aligned} \tag{3.4}$$

where the mean-variance relation has been used to write the covariance matrix as a function of λ and ϕ, and we use the notation

$$\boldsymbol{\Sigma}' = \mathrm{diag}(\boldsymbol{\lambda}^c).$$

Cao et al. propose a time varying method for network tomography, where they take a short window of measurements, and use the Expectation Maximization (EM) algorithm [MK97] to find the maximum likelihood estimate of the traffic matrix for that time slot. Once a time series of estimates is obtained, additional smoothing is performed on those values.

The problem with the time-varying network tomography is that it is rather time consuming. The authors conclude that it does not readily scale for realistic size networks. In [CWYZ01] Cao et al. present a way to modify their method in [CDWY00] in order to use it for realistic size networks. They call the new method a divide-and-conquer approach, as the idea is to form several smaller subproblems and solve them separately.

The Pseudo likelihood approach [LY03] is another scalable method for likelihood estimation. The idea is to use a slightly modified EM algorithm. The problem is divided into subproblems with each covering one OD-pair. The solution of this calls for the use of Multiple-step Gradient EM algorithm. The computational complexity of each EM step is now $O(n^{3.5})$ compared to the full likelihood method's $O(n^5)$. The authors report that for a small network studied the error in estimation accuracy increases only from 8% to 9% when switching from the full likelihood to the pseudo likelihood method.

Second Generation Methods Using Spatial Correlation

This class of estimation methods use a model where the extra information comes from a spatial model that describes the relationship between the OD flows. This is most commonly the gravity model or some other form of description for the nodes' fanout.

The gravity model is named after Newton's law of gravitation. As in the law of gravitation the force between two objects is proportional to the masses of the objects and the inverse of the square of the distance between them.

Similarly in the gravity modelling for data networks the traffic between two nodes is assumed to be proportional to the total traffic volumes of those nodes. Gravity models have been used in social science to model the movement of people or goods between two areas, as well as in telephone networks [LaJW97, Pyh63, Tin62]. The idea is that if we have no knowledge of where a bit is coming from or where it is going to, the best guess is to make the estimate proportional to traffic volumes sent and received by each node in the network.

The general form of the gravity model has a repulsion term and an attraction term that are multiplied together and then divided by a distance function. In the case of traffic matrix estimation it can be written in the form proposed by Kowalski and Warfield [KW95] for teletraffic demands:

$$X_{sd} = k_s \frac{O_s T_d}{d_{sd}^{\alpha_s}}. \tag{3.5}$$

The repulsion term is O_s which is the total traffic originating from node s. The attraction term is T_d, the total traffic terminating at node d. The numerator d_{sd} is a distance function between nodes s and d, where α_s is the distance parameter. Coefficient k_s is a normalizing constant.

Zhang et al. [ZRDG03] use an approach where the normalizing coefficient and distance function are put together to form a friction factor between origin and destination. However, they notice that the inference of the friction factors is an equivalent problem to the traffic matrix inference, and thus ill-posed. Hence they simplify their model to

$$X_{sd} = k \cdot T_s^{in} T_d^{out}, \tag{3.6}$$

where $T_s^{in/out}$ is the amount of traffic entering/leaving the network through node s and the normalizing constant is

$$k = \sum_i T_i^{out} \qquad \text{or} \qquad k = \sum_i T_i^{in},$$

with both yielding identical results.

Zhang et al. [ZRDG03] also generalize the gravity model to handle additional information, specifically to use the knowledge that some of the egress links are peering links to other ISPs while others are access links, to differentiate between customer and peering traffic. They use this method to obtain a starting point $\boldsymbol{x}_0$, and solve the quadratic programming problem of the L_2 norm of a vector, e.g., the euclidian distance

$$\begin{array}{rl} \text{min} & ||(\boldsymbol{x} - \boldsymbol{x}_0)/\boldsymbol{w}|| \\ \text{so that} & ||\boldsymbol{A}\boldsymbol{x} - \boldsymbol{y}|| = 0, \end{array} \tag{3.7}$$

where $\boldsymbol{w}$ is a weight vector. This tomogravity approach utilizes the link count information.

In [ZRLD03] the tomogravity method is generalized using an information theoretic approach. The gravity model is based on independence between origin and destination of the traffic. In information theoretic terms the independence between source and destination, implied by the gravity model, is equivalent to the mutual information being zero. As the mutual information is also always positive, it is thus an appropriate penalty function to be used in a regularized minimization problem.

$$\begin{array}{rl} \min\limits_{\boldsymbol{x}} & ||\boldsymbol{y} - \boldsymbol{A}\boldsymbol{x}||^2 + \lambda^2 \sum\limits_{i:\, g_i > 0} \frac{x_i}{N} \log\left(\frac{x_i/N}{g_i}\right) \\ \text{subject to} & x_i \geq 0, \end{array} \tag{3.8}$$

where g_i is an element of $\boldsymbol{g}$, the gravity model estimate.

That is, we want a solution that is a tradeoff between satisfying the link count relation and having an a priori plausibility, which here means that the mutual information is small and the solution is thus close to the gravity model.

Medina at al. [MTS$^+$02, MST$^+$03, MST$^+$04] introduce the choice model for POP to POP traffic matrix estimation, where they combine the attraction term of the destination node with the distance function to form a fanout term α_{sd} that determines which portion of the traffic from a source node s is going to each destination node d. The choice model is thus written as

$$X_{sd} = O_s \alpha_{sd}. \tag{3.9}$$

The authors use a Discrete Choice Model (DCM) to estimate the fanouts.

Gunnar et al. [GJT04] analyze real traffic from the Global Crossing backbone network. They find that the fanout factors remain constant over time, even while the traffic amounts fluctuate due to a diurnal pattern. The fanout α_{sd} gives the percentage of traffic that source node s sends to destination node d of its total traffic. Thus the fanouts sum to unity for each source node. Assuming the constant fanouts they write

$$\boldsymbol{x}_t = S_t \boldsymbol{\alpha},$$

where S_t is a time dependent scaling term and $\boldsymbol{\alpha}$ is the vector of fanout terms, which sum to unity for each source node. The link count equations get the form

$$\begin{aligned} \boldsymbol{y}_t &= \boldsymbol{A} S_t \boldsymbol{\alpha}, \\ \sum_d \alpha_{sd} &= 1 \qquad \forall s. \end{aligned} \tag{3.10}$$

With a time series of link counts, the above system will quickly become overdetermined, and there will be a unique solution vector $\boldsymbol{\alpha}$.

Third Generation Methods Using Direct Measurements

Soule et al. [SLT+05] propose a principal component method that makes use of the eigenflow representation by Lakhina et al. [LPC+04]. It is a spatial model making use of direct Netflow measurements. While the traffic matrix inference is ill-posed problem, the dimensionality of the aforementioned components is so much smaller that they can be estimated from the link counts, and thus the traffic matrix can be estimated as well, as long as the method is first calibrated by direct measurements.

Papagianniki et al. [PTL04] do not use the routing matrix, but rely on measurements alone to obtain the traffic matrix. In their spatio-temporal model the fanouts for each node are defined from calibration done by direct measurements such as Netflow. These are then combined with SNMP measurements to obtain an estimate. The nodes check periodically the validity of the fanouts by direct measurements and perform a new calibration if necessary.

The Kalman Filtering method by Soule et al. [SLT+05, SSNT05] is another approach using a spatio-temporal model. It also requires calibration by direct measurement and the uses Kalman filter to model the traffic evolution according to a linear system capturing both the temporal evolution of an OD pair and cross-correlation between OD pairs, should any exist. Re-calibration is performed whenever the theoretical link counts based on the estimates deviate too much from the observed link counts.

Liang et al. [LTY06] propose a method, where calibration measurements are done on one link only in each estimation interval. This method is shown to provide accurate estimates with low measurement overhead.

Other Methods

In the Bayesian approach the idea is to compute conditional probability distribution for OD pair traffic, given the link counts and a prior distribution. While gravity estimate is the logical candidate for the prior distribution, this is purely a computational technique and thus difficult to classify in the above classes.

Tebaldi and West [TW98] use a Poisson-distribution, and the joint distribution of OD pair counts and their expectations is conditioned on the observed link counts. Analytical computations are difficult in this case, and thus Markov Chain Monte Carlo (see e.g. [RC99]) methods are used to obtain the posterior distribution.

Vaton and Gravey [VG03] make use of several successive link count measurements in their iterative Bayesian method that allows for modulated process for the underlying traffic matrix distribution. The method consists of iteration and exchange of information between two ”boxes” as depicted in Figure 3.1. The first box follows the method by Tebaldi and West, and

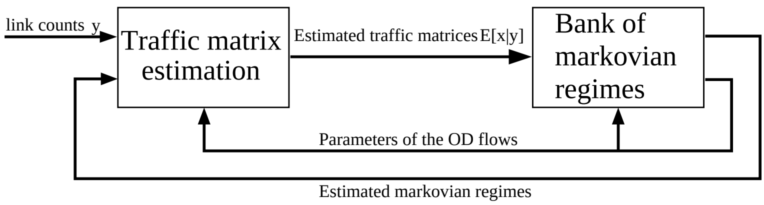

Figure 3.1: The Vaton-Gravey iterative method

simulates the traffic matrix from the link counts at each fixed time period using MCMC methods, and utilizing some prior distribution. The estimated traffic matrices are then given to the second box, where the values are fitted to the underlying model and the parameters for that model are estimated using maximum likelihood estimation.

There have been some attempts to solve the problem by linear programming methods [Gol00, EMH05, CL03]. These approaches work well on some small examples, but not necessarily for realistic size networks. Rahman et al. [RSCA06] point out that while Conway and Li [CL03] achieve good performance, the level graphs they use are not representative of any realistic network.

The problem with these approaches is that any point on the plane $\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x}$ gives the same value for the objective functions, as the whole feasible region is pareto-optimal. So while the real answer does yield the maximum value for the objective function with these weights, so would any other feasible answer. We have not gained any new knowledge about the situation by formulating it as an LP problem.

Vaton et al. [VBG05] note that a classical method to solve underdetermined linear systems is to minimize the euclidian norm. However, they conclude that this is not a realistic approach for the traffic matrix estimation problem, as it finds the solution that has the OD pairs as close to same size with each other as possible, which is not a realistic criterion.

A reasonable use for LP is given in [GJT04], where the authors formulate worst case bounds for OD counts. That is, they use linear programming to find bounds for possible values of OD counts. Obvious bounds are zero for lower bound and the lowest link count on the OD pair's path for upper bound, but in many cases it is possible to find tighter bounds for some OD pairs. This method, however, is quite heavy computationally, as two LP problems need to be solved for each OD pair.

Soule et al. [SNC+04] propose a method that achieves accurate results by changing routing. The idea is that in different routing scenarios different OD pairs are easier to estimate, as the links they use might be less heavily populated by other flows in some routing schemes than others. Each routing scheme is used for a period of multiple measurements, and several different routings are used. The extra information needed in underconstrained problems comes from the additional routing scenarios. The problem of designing the weight changes to optimally generate the different routings is tackled by the authors in [NCTD04].

3.5 Contribution of the Thesis

We study the traffic matrix estimation problem in Publications 4, 5, and 6. The work consists of evaluation of methods, statistical calculation of performance bounds of the likelihood method as well as two novel estimation methods.

Publication 4 derives the statistical Cramér-Rao lower bounds (CRLB) for the maximum likelihood approach of traffic matrix estimation. We derive analytically the Fisher information matrix under this framework and obtain the Cramér-Rao lower bound for the variance of an estimator of the traffic matrix. Applications for the use of the CRLB are then demonstrated. From the bounds we can directly obtain confidence intervals for maximum likelihood estimates, which can be used to evaluate the efficiency of an estimator, or to find an optimal location for direct measurements.

The proposed second moment methods in literature are computationally very heavy and thus time consuming. A quicker way to obtain an estimate using the mean variance relation was missing. Thus, in Publication 5, a light-weight second moment method is proposed. It makes use of the fact that not only is the system of first and second moment link measurement statistics identifiable, but we can solve the OD-pair variances using only the link count covariance matrix.

Publication 6 studies the effect of the two sources of extra information on the estimate through a simulation study, evaluating the accuracy of the estimates when the extra information is less than completely accurate. In addition, a novel estimation technique is proposed, which incorporates both sources of extra information.

3.6 Cramér-Rao Bounds for Maximum Likelihood Method

In general, we can say that there is a trade-off between the computational complexity and the accuracy of the estimate. However, no matter how elaborate the technique, there is a bound for the accuracy of the estimate.

This is due to the stochastic nature of the traffic process, which makes it impossible to obtain estimate accuracy better than a certain level.

The traffic volume can be considered a random variable. The Fisher information matrix gives the amount of information that the observed traffic volumes carry of the underlying parameter, namely the expected traffic volume. For any unbiased estimate, the Cramér-Rao lower bound (CRLB), which is the inverse of the Fisher information matrix, gives the limit of how small variances it is possible to obtain for an estimator [Sha98]. In Publication 4 we calculate the Cramér-Rao bounds for the traffic matrix estimation problem.

Cramér-Rao Lower Bound

In our model we assume that the OD-pair traffic follows Gaussian distribution, that OD pairs are independent of each other, and also that successive measurements for each OD pair are independently and identically distributed and that the functional mean-variance relation holds. The link counts have probability density function (pdf) $p(\boldsymbol{y}; \boldsymbol{\Psi})$, where $\boldsymbol{\Psi}$ is the vector containing the unknown parameters.

Theorem 3.6.1 Under regularity conditions the covariance matrix of any unbiased estimator $\boldsymbol{\Psi}^*$ satisfies

$$C_{\boldsymbol{\Psi}^*} - \mathcal{I}^{-1}(\boldsymbol{\Psi}) \geq 0, \tag{3.11}$$

where "≥ 0" is interpreted so that the matrix is positive semidefinite, and $\mathcal{I}(\boldsymbol{\Psi})$ is the Fisher information matrix evaluated at the true value of $\boldsymbol{\Psi}$.

The above theorem gives the Cramér-Rao lower bound. It states that $C_{\boldsymbol{\Psi}^*}$, the variance/covariance matrix of any unbiased estimator cannot be lower than the inverse of the Fisher information matrix.

The incomplete data is a multivariate Gaussian with mean $\boldsymbol{\mu} = \boldsymbol{\mu}(\boldsymbol{\Psi})$ and covariance matrix $\boldsymbol{C} = \boldsymbol{C}(\boldsymbol{\Psi})$. The probability density function is

$$\begin{aligned} p(\boldsymbol{y}; \boldsymbol{\Psi}) &= \frac{1}{(2\pi)^{m/2} \det \boldsymbol{C}(\boldsymbol{\Psi})^{1/2}} \cdot \\ & \cdot \exp\left\{ -\frac{1}{2} (\boldsymbol{y} - \boldsymbol{\mu}(\boldsymbol{\Psi}))^{\mathrm{T}} \boldsymbol{C}(\boldsymbol{\Psi})^{-1} (\boldsymbol{y} - \boldsymbol{\mu}(\boldsymbol{\Psi})) \right\} \end{aligned} \tag{3.12}$$

It follows that the log-likelihood is

$$\begin{aligned} l(\boldsymbol{y}; \boldsymbol{\Psi}) = & -\log(2\pi)^{m/2} \\ & -\frac{1}{2} \log \det(\boldsymbol{C}(\boldsymbol{\Psi})) - \frac{1}{2} (\boldsymbol{y} - \boldsymbol{\mu}(\boldsymbol{\Psi}))^{\mathrm{T}} \boldsymbol{C}(\boldsymbol{\Psi})^{-1} (\boldsymbol{y} - \boldsymbol{\mu}(\boldsymbol{\Psi})). \end{aligned} \tag{3.13}$$

An element of the information matrix can be written as

$$\mathcal{I}(\boldsymbol{\Psi})_{ij} = E\left[\frac{\partial l(\boldsymbol{y}; \boldsymbol{\Psi})}{\partial \Psi_i} \frac{\partial l(\boldsymbol{y}; \boldsymbol{\Psi})}{\partial \Psi_j} \right]. \tag{3.14}$$

The following analytical expression for the information matrix is derived in the appendix of Publication 4

$$\mathcal{I}(\boldsymbol{\Psi})_{ij} = \frac{\partial \boldsymbol{\mu}(\boldsymbol{\Psi})^{\mathrm{T}}}{\partial \Psi_i} \boldsymbol{C}^{-1}(\boldsymbol{\Psi}) \frac{\partial \boldsymbol{\mu}(\boldsymbol{\Psi})}{\partial \Psi_j} + \frac{1}{2} tr \left(\boldsymbol{C}^{-1}(\boldsymbol{\Psi}) \frac{\partial \boldsymbol{C}(\boldsymbol{\Psi})}{\partial \Psi_i} \boldsymbol{C}^{-1}(\boldsymbol{\Psi}) \frac{\partial \boldsymbol{C}(\boldsymbol{\Psi})}{\partial \Psi_j} \right). \tag{3.15}$$

For the traffic matrix estimation problem the link counts $\boldsymbol{y}_t$ obey a multivariate Gaussian distribution with mean

$$\boldsymbol{\mu}(\boldsymbol{\Psi}) = \boldsymbol{A}\boldsymbol{\lambda}, \tag{3.16}$$

and covariance matrix

$$\boldsymbol{C}(\boldsymbol{\Psi}) = \phi \boldsymbol{A} \boldsymbol{\Sigma}' \boldsymbol{A}^{\mathrm{T}}, \tag{3.17}$$

where ϕ is a scaling parameter.

Using these, we can calculate the information matrix for the traffic matrix estimation problem. The information matrix has the following structure

$$\mathcal{I}_t(\boldsymbol{\Psi}) = \left(\begin{array}{c|c} I_1 & I_2 \\ \hline I_3 & I_4 \end{array} \right), \tag{3.18}$$

where I_1 is a $n \times n$ matrix, I_2 is a column vector of length n, I_3 is a row vector of the same length, and I_4 is a scalar. To simplify the notation we introduce the matrix

$$\boldsymbol{W} = \boldsymbol{A}^{\mathrm{T}} (\boldsymbol{A} \boldsymbol{\Sigma} \boldsymbol{A}^{\mathrm{T}})^{-1} \boldsymbol{A}, \tag{3.19}$$

which has the elements

$$w_{ij} = {\boldsymbol{A}^{j}}^{\mathrm{T}} (\boldsymbol{A} \boldsymbol{\Sigma} \boldsymbol{A}^{\mathrm{T}})^{-1} \boldsymbol{A}^{i}. \tag{3.20}$$

Now the elements of the matrix are given as

$$(I_1)_{i,j} = w_{ij} + \frac{c^2 \lambda_i^{c-1} \lambda_j^{c-1}}{2} \phi^2 w_{ij}^2. \tag{3.21}$$

For $i = 1, \cdots, n$

$$(I_2)_{i,(n+1)} = \frac{c \lambda_i^{c-1}}{2} w_{ii}. \tag{3.22}$$

Analogously, for $j = 1, \cdots, n$

$$(I_3)_{(n+1),j} = \frac{c \lambda_j^{c-1}}{2} w_{jj}. \tag{3.23}$$

And finally,

$$(I_4)_{(n+1),(n+1)} = \frac{m}{2\phi^2}. \tag{3.24}$$

We have now obtained an analytical expression for the Fisher information matrix of the traffic matrix estimation problem. The Cramér-Rao lower bound for the variance of an estimator is then just $\mathcal{I}^{-1}$, where the CRLB for variances of the parameters are the diagonal elements.

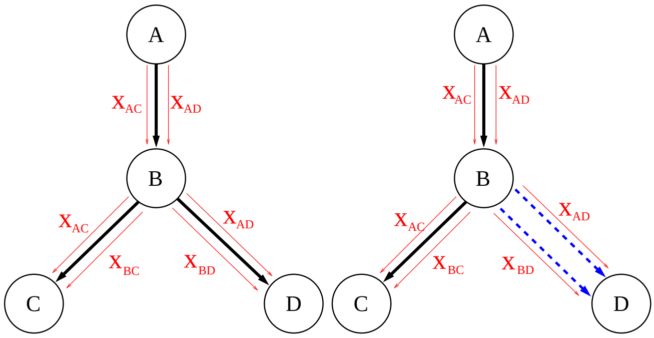

Figure 3.2: Left: Example topology. Right: Link BD is replaced by virtual links BD_1 and BD_2

Applications of CRLB

The asymptotic efficiency of the MLE is a well-known result [MK97]. It follows that the asymptotic covariance matrix of the MLE is equal to the inverse of the expected information matrix, that is, the CRLB.

Thus, when evaluating the performance of an estimation method in simulation studies with synthetic traffic matrices, we can obtain sample variances for the considered methods. Then calculating the Cramér-Rao bounds it is possible to compare the sample variances to the bounds, and thus to the variance of MLE. This way we can evaluate how much less accurate the methods are than the full MLE, without having to run the full likelihood method.

Another application is to use the bounds to analytically assess which links would be most beneficial for a direct measurement location in the sense of minimizing the uncertainty in the traffic matrix estimate. To incorporate the direct measurements of some OD flows to the traffic matrix estimation framework we propose a model that creates a new linear system. This can be interpreted as a virtual topology, where the link on which the direct measurements are made is replaced by several virtual links, such that each OD pair using the link would have its own virtual link, see Figure 3.2. This enables us to incorporate the direct OD-pair measurements without changing the basic formulation of the estimation problem. The Cramér-Rao bounds are calculated for the new system to find out the variance, and thus expected error, of the OD pairs. Comparing the average errors of different situations enables us to select the measurement location that decreases the error most.

A third application is to use the bounds to form a confidence interval polytope for the estimate and use this in robust load balancing. This approach is studied in more detail in section 4.5.

3.7 The Quick Method Based on Link Covariances

In Publication 5 we propose the quick method for traffic matrix estimation, which uses the link count covariances as the extra information. Maximum likelihood estimation (MLE) uses the second moment statistic, the link count covariance, as the additional information that is needed to yield an estimate. It is also necessary to assume local stationarity for the measurements considered, and a distribution which the stochastic fluctuation of the traffic follows.

The MLE relies on the fact that the system of first and second order link count statistics together make the system identifiable with regard to the first order OD-pair statistics, i.e. we are able to find solution for the likelihood equations if there exists a functional relationship between the mean and the variance of OD-pair traffic. As discussed in section 2.6, the commonly used relation is the power-law relation

$$\boldsymbol{\Sigma} = \phi \cdot \operatorname{diag}\{\boldsymbol{\lambda}^c\}. \tag{3.25}$$

But, in fact, the second order statistic for OD-pairs is identifiable based solely on the second order statistics of the link counts, as long as we assume independence among OD-pairs and a sensible routing scheme. This result is proven by Soule et al. [SNC$^+$04]. Since we can analytically solve the variance of the OD-pairs by the least square method, and the power-law relation between variance and mean is assumed, we can then solve the traffic matrix from our variance estimate.

Solving OD-pair covariance matrix from link counts

Let us denote the number of links by J and the number of OD-pairs by N. Then the vector form of traffic matrix $\boldsymbol{x}$ has the dimension $(N \times 1)$, link loads $\boldsymbol{y}$ has the dimension $(J \times 1)$.

First, let us define $\boldsymbol{S}^{(y)}$ as a $\frac{1}{2}J(J+1)$-vector containing diagonal and upper triangle elements of the link covariance matrix $\boldsymbol{\Sigma}^{(y)}$. Define $\boldsymbol{S}^{(x)}$ as an N-vector containing the diagonal elements of the OD-pair covariance matrix $\boldsymbol{\Sigma}^{(x)}$. $\boldsymbol{A}$ is the routing matrix. Then define a $(\frac{1}{2}J(J+1) \times N)$ matrix $\boldsymbol{B}$ that relates vector $\boldsymbol{S}^{(y)}$ to vector $\boldsymbol{S}^{(x)}$. A row of $\boldsymbol{B}$ is indexed by a compound index (ij) where $i = 1, \ldots, J;\ j = i, \ldots, J$, meaning that the index runs through $\frac{1}{2}J(J+1)$ values,

$$B_{(ij),k} = A_{i,k}A_{j,k} \qquad \begin{array}{l} i = 1, \ldots, J;\ j = i, \ldots, J \\ k = 1, \ldots, N. \end{array}$$

The rows of $\boldsymbol{B}$ indicate the elements of $\boldsymbol{x}$ contributing to covariance between links i and j. In vector notation, we can write the relation as

$$\boldsymbol{S}^{(y)} = \boldsymbol{B}\boldsymbol{S}^{(x)}. \tag{3.26}$$

Typically $\frac{1}{2}J(J+1) > N$ and equation (3.26) is overdetermined. The least square estimate (LSE) solution (see, e.g., [Lue69]), to the equation is

$$\boldsymbol{S}^{(x)} = (\boldsymbol{B}^{\mathrm{T}}\boldsymbol{B})^{-1}\boldsymbol{B}^{\mathrm{T}}\boldsymbol{S}^{(y)}. \tag{3.27}$$

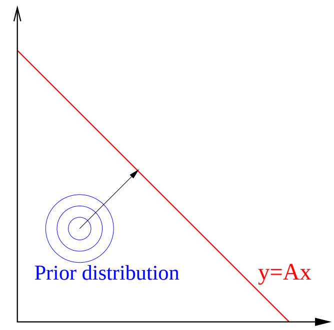

Figure 3.3: Projection from prior distribution starting point to link constraint plain

Projection method

Now that we have an estimate for the variances of each OD-pair, it is trivial to find an estimate of the mean by using the mean-variance relation (3.25).

$$\boldsymbol{\lambda}_0 = (\phi^{-1}\boldsymbol{S}^{(x)})^{\frac{1}{c}}. \tag{3.28}$$

The problem with this estimate is that it does not require the solution to satisfy the link count equation (3.3), which is a stronger condition than the second moment relation. The preliminary estimate $\boldsymbol{\lambda}_0$ can be improved by projecting the result to the surface that satisfies the first moment condition. This yields our estimate

$$\boldsymbol{\lambda} = \boldsymbol{\lambda}_0 + \boldsymbol{A}^{\mathrm{T}}(\boldsymbol{A}\boldsymbol{A}^{\mathrm{T}})^{-1}(\boldsymbol{y} - \boldsymbol{A}\boldsymbol{\lambda}_0). \tag{3.29}$$

This is an unweighted projection, as depicted in Figure 3.3, in the sense that uncertainty of the prior estimate is the same in all directions.

The projection might yield negative values for small OD pairs, as no positivity constraint is imposed. In order to keep the method as light-weight as possible, we simply substitute the negative estimates by zero, concluding that these OD pairs are negligibly small.

Constrained minimization

Another approach, developed in Publication 5, is to require the condition $\boldsymbol{y} = \boldsymbol{A}\boldsymbol{\lambda}$ to be satisfied from the outset, and try to satisfy the mean-variance relation in the least square sense. We get a constrained minimization problem

$$\begin{aligned} \min_{\boldsymbol{\lambda},\phi} \quad & \|\boldsymbol{S}^{(y)} - \boldsymbol{B}\phi\boldsymbol{\lambda}^c\| \\ \text{subject to} \quad & \boldsymbol{y} = \boldsymbol{A}\boldsymbol{\lambda}. \end{aligned} \tag{3.30}$$

In general, this has to be solved numerically. However, in the special case of $c = 1$ an explicit solution can be derived.

Introducing a vector of Lagrange multipliers $\boldsymbol{\alpha}$, the objective function to be minimized can be written as

$$\begin{aligned} f(\boldsymbol{\lambda}, \boldsymbol{\alpha}, \phi) &= (\boldsymbol{S}^{(y)} - \phi \boldsymbol{B}\boldsymbol{\lambda})^{\mathrm{T}}(\boldsymbol{S}^{(y)} - \phi \boldsymbol{B}\boldsymbol{\lambda}) + 2\boldsymbol{\alpha}^{\mathrm{T}}(\overline{\boldsymbol{y}} - \boldsymbol{A}\boldsymbol{\lambda}) \\ &= \phi^2 \boldsymbol{\lambda}^{\mathrm{T}} \boldsymbol{B}^{\mathrm{T}} \boldsymbol{B}\boldsymbol{\lambda} - 2\phi {\boldsymbol{S}^{(y)}}^{\mathrm{T}} \boldsymbol{B}\boldsymbol{\lambda} - 2\boldsymbol{\alpha}^{\mathrm{T}} \boldsymbol{A}\boldsymbol{\lambda} \\ &\quad + {\boldsymbol{S}^{(y)}}^{\mathrm{T}} \boldsymbol{S}^{(y)} + 2\boldsymbol{\alpha}^{\mathrm{T}} \overline{\boldsymbol{y}}. \end{aligned} \tag{3.31}$$

The above expression is quadratic in $\boldsymbol{\lambda}$, and the minimum with respect to $\boldsymbol{\lambda}$ can easily be found,

$$\boldsymbol{\lambda} = \phi^{-2}(\boldsymbol{B}^{\mathrm{T}}\boldsymbol{B})^{-1}(\boldsymbol{A}^{\mathrm{T}}\boldsymbol{\alpha} + \phi \boldsymbol{B}^{\mathrm{T}}\boldsymbol{S}^{(y)}). \tag{3.32}$$

The Lagrange multipliers $\boldsymbol{\alpha}$ are then determined such that the constraints are satisfied:

$$\overline{\boldsymbol{y}} = \boldsymbol{A}\phi^{-2}(\boldsymbol{B}^{\mathrm{T}}\boldsymbol{B})^{-1}(\boldsymbol{A}^{\mathrm{T}}\boldsymbol{\alpha} + \phi \boldsymbol{B}^{\mathrm{T}}\boldsymbol{S}^{(y)}), \tag{3.33}$$

from which

$$\boldsymbol{\alpha} = (\phi^{-2}\boldsymbol{A}(\boldsymbol{B}^{\mathrm{T}}\boldsymbol{B})^{-1}\boldsymbol{A}^{\mathrm{T}})^{-1}(\overline{\boldsymbol{y}} - \phi^{-1}\boldsymbol{A}(\boldsymbol{B}^{\mathrm{T}}\boldsymbol{B})^{-1}\boldsymbol{B}^{\mathrm{T}}\boldsymbol{S}^{(y)}). \tag{3.34}$$

Minimizing $f(\boldsymbol{\lambda}, \boldsymbol{\alpha}, \phi)$ with respect to ϕ yields

$$\phi = (\boldsymbol{\lambda}^{\mathrm{T}}\boldsymbol{B}^{\mathrm{T}}\boldsymbol{B}\boldsymbol{\lambda})^{-1}{\boldsymbol{S}^{(y)}}^{\mathrm{T}}\boldsymbol{B}\boldsymbol{\lambda}. \tag{3.35}$$

Substitution of (3.34) into (3.32) gives $\boldsymbol{\lambda}$ as a function of ϕ

$$\boldsymbol{\lambda} = \boldsymbol{K}\overline{\boldsymbol{y}} - \phi^{-1}\big(\boldsymbol{K}\boldsymbol{A}(\boldsymbol{B}^{T}\boldsymbol{B})^{-1}\boldsymbol{B}^{T}\boldsymbol{S}^{(y)} + \boldsymbol{B}^{T}\boldsymbol{S}^{(y)}\big),$$

where we use the notation

$$\boldsymbol{K} = (\boldsymbol{B}^{T}\boldsymbol{B})^{-1}\boldsymbol{A}^{T}(\boldsymbol{A}(\boldsymbol{B}^{T}\boldsymbol{B})^{-1}\boldsymbol{A}^{T})^{-1}.$$

Substituting $\boldsymbol{\lambda}$ further in (3.35) yields an quadratic equation for ϕ, which is easily solvable. This solution can be then substituted back to (3.34) and (3.32) to obtain the explicit expression for $\boldsymbol{\lambda}$.

Simulation Results

We use a 12-node topology and synthetic traffic generated with the parameter value $c = 1.5$. The traffic volumes for the OD pairs vary so that the largest are approximately hundred times as large as the smallest ones. This creates great difficulties for the quick methods regarding the estimation of the smaller OD pairs. The estimates of the projection method for the smallest OD pairs are far off the real traffic volumes. Due to the fact that the estimates for some of the smallest OD pairs have errors of several hundred percent, the mean relative error is also affected greatly by these, and is 59% for the projection method and 110% for the constrained optimization, while it is 29% for the MLE.

However, the most important thing is to estimate the largest OD pairs. If we concentrate only on the largest OD pairs that comprise 90% of total traffic in volume, the projection method is more competitive. The errors are 27% for the projection method and 19% for the MLE.

3.8 The Combined Method

In Publication 6 we propose an estimation method combining two sources of extra information. As stated before, the traffic matrix estimation problem is underconstrained and some extra information has to be brought into the situation to get a unique estimate. The accuracy of this estimate depends on the relevance of the extra information, viz. the validity of the assumptions made in order to use the information in the estimation. Current methods utilize either gravity model or the mean-variance relation. However, as both are relevant information to the problem, we propose the combined method that utilizes both.

There are two ways to take both sources of extra information into account. We can write both starting points into the regularization equation, and optimize them simultaneously. The objective function becomes

$$\min\left\{\|\boldsymbol{y}-\boldsymbol{A}\boldsymbol{x}\|^2+\lambda^2\sum\frac{x_i}{N}\log\left(\frac{x_i}{g_i}\right)+\mu^2\sum\frac{x_i}{N}\log\left(\frac{x_i}{q_i}\right)\right\}, \tag{3.36}$$

where $\boldsymbol{g}$ is the gravity model prior and $\boldsymbol{q}$ is the quick method prior of (3.28).

Another possibility is to just take a componentwise average of the two priors and insert the resulting combined prior into the regularization function. This yields

$$\min\left\{\|\boldsymbol{y}-\boldsymbol{A}\boldsymbol{x}\|^2+\lambda^2\sum\frac{x_i}{N}\log\left(\frac{x_i}{wg_i+(1-w)q_i}\right)\right\} \tag{3.37}$$

as the objective function.

It turns out that the latter method, which is computationally simpler, also outperforms the first method. Thus we concentrate on that approach.

3.9 Sensitivity of Methods to Their Underlying Assumptions

The common sources of extra information in traffic matrix estimation are the gravity model assumption and the mean-variance relation. Above we have introduced methods based on these, as well as the combined method using both sources of information. If these assumptions are inaccurate the accuracy of the methods obviously suffers. In Publication 6 we study the effect that inaccuracies in these underlying assumptions have on the gravity model method, second moment method and combined method, respectively, to find out how much the inaccuracies in the assumptions affect the accuracy of the traffic matrix estimates.

To remove the effect of different estimation techniques we use for each method the regularized minimization problem with a penalty function. In this case

$$\min_{\boldsymbol{x}}\left[||\boldsymbol{y}-\boldsymbol{A}\boldsymbol{x}||_2^2+\lambda^2 J(\boldsymbol{x})\right], \tag{3.38}$$

where λ is a regularization parameter and J is a penalization functional.

The penalty function is given as

$$J(\boldsymbol{x}) = \sum_i \frac{x_i}{N} \log\left(\frac{x_i}{f_i}\right),$$

where f_i is the quick prior, the gravity prior or a weighted average of these as in (3.37).

Simulation Methodology

In the simulation study we create synthetic data sets in which the assumptions are true to various degrees, starting from a perfect fit and then making it gradually worse. For any given situation, the goodness of fit value of the mean-variance relation is placed on the vertical axis and the goodness of fit of the gravity model on the horizontal axis of a diagram. Then, at each point of this grid we can make a simulation study to find out which method is more accurate with these particular goodness of fit values for the assumptions, and eventually find out where do the equivalence curves lie, which indicate that the methods are equally accurate at those points.

We use a 12-node backbone topology and generate OD-pair traffic volumes following the gravity model. For each simulation we draw identically and independently distributed samples from a Gaussian distribution with parameter vector $\boldsymbol{\lambda}$ for the means of the OD pairs and covariance matrix $\boldsymbol{\Sigma}$ defining the variances of the OD pairs. To achieve data sets which follow the gravity model assumption and the mean-variance relation only to some degree we add a random component to these parameters.

The deterministic component for the mean is denoted by $\boldsymbol{\lambda}_g$, and follows the gravity model exactly. We add to this a random component

$$\boldsymbol{\epsilon}_\lambda \in (-\lambda_g, \lambda_g),$$

and produce synthetic data with a desired goodness of fit value with regard to the gravity model by changing the weight $w \in [0, 1]$ of the error term

$$\boldsymbol{\lambda} = \boldsymbol{\lambda}_g + w\boldsymbol{\epsilon}_\lambda.$$

For the variance, the final parameter value is taken similarly with a deterministic term following the mean-variance relation and an error term multiplied by a weight coefficient

$$\sigma^2 = \sigma_m^2 + v\epsilon_\sigma.$$

To produce the synthetic data set we draw T samples of traffic counts $\boldsymbol{x}_t$ from a Gaussian distribution with the parameter values obtained above.

$$\boldsymbol{x}_t \sim \mathrm{N}(\boldsymbol{\lambda}, \boldsymbol{\Sigma}).$$

To measure how close the synthetic OD pair parameters are to the deterministic values we use the goodness of fit

$$R^2 = 1 - \frac{\mathrm{ESS}}{\mathrm{TSS}},$$

where ESS is the error sum of squares and TSS the total sum of squares between the actual parameter values and the deterministic parameter values. If the fit is perfect, then ESS is zero and $R^2 = 1$. If the traffic matrix is totally random, then ESS is approximately same as TSS and $R^2 \approx 0$.

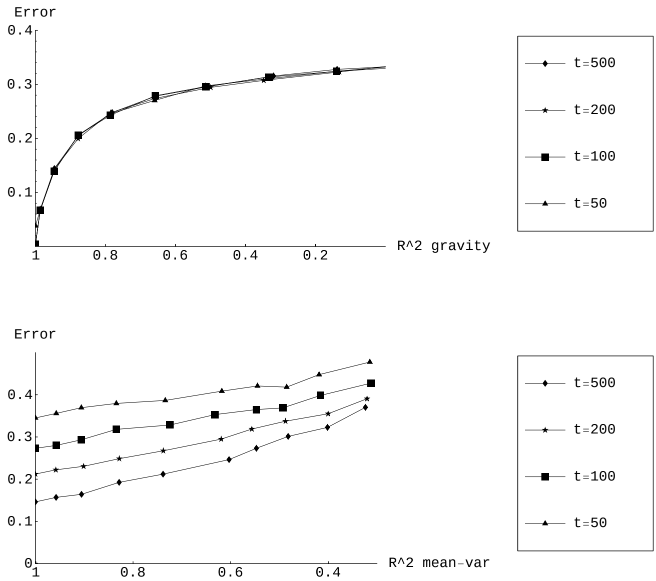

Figure 3.4: Effect on estimation accuracy when assumptions goodness of fit deteriorates. Top: Gravity model assumption. Bottom: Mean-variance relation.

Simulation Results

We use the mean relative error of the largest OD pairs, which comprise 80% of total traffic. On the top of Figure 3.4 the errors of the gravity model estimate are displayed as a function of the R^2 value of the goodness of fit of the gravity assumption. On the bottom of Figure 3.4 a similar situation is shown regarding the second moment method. For the gravity model the sample size does not affect the accuracy. They are also very accurate when the gravity model assumption holds but quickly grow worse when the fit becomes less exact. On the other hand, the second moment method depends significantly on the sample size, due to the fact that it is difficult to get accurate estimates of the sample variance with small sample sizes. The accuracy of the estimator does not deteriorate very quickly with a decreasing goodness of fit of the mean-variance relation.

It is to be expected that if the gravity model holds well, while the mean-variance relation does not, the gravity based methods are more accurate, and vice versa. In Figure 3.5 equivalence curves are displayed, indicating when the two estimators compared are equally accurate. On the left side of the diagrams near the vertical axis the gravity assumption holds exactly, and on the bottom of the diagram near the horizontal axis the mean-variance relation holds exactly. Thus, on the area from the equivalence curve to the top left corner the gravity model method is more effective and from the curve to the bottom right corner the second moment methods yield better results. As the latter method was dependent on the sample size, also the

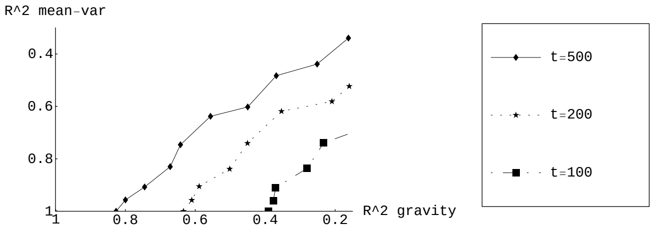

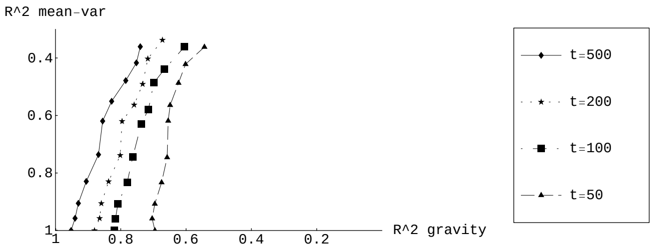

Figure 3.5: Equivalence curves showing which R^2 values of gravity yield similar estimation errors. Top: Gravity vs. 2nd moment method. Bottom: Gravity vs Combined method.

Table 3.1: R^2 values calculated from real data sets

	Mean-var relation	Gravity model
Abilene	0.76	0.84
Funet	0.83	N/A
Lucent	0.76	0.96

equivalence curves depend on the sample size.

Typical R^2 values for real data sets (Funet, Lucent[CDWY00] and Abilene[1]) available are listed in Table 3.1. We notice that these values fall onto the area which is to the left of the equivalence curves. Thus, even with a very large sample size of 500 the gravity model seems to be the better choice for estimating the traffic matrix in these networks. This strongly indicates that the gravity model is the better of the two approaches.

Equivalence curves comparing the combined method and the gravity method are depicted on the bottom of Figure 3.5. Again, the gravity model is the best one to the left of the equivalence curves and the combined method is the best on the right side of the curves. We notice that the combined method is rather accurate, and the fit of the gravity model assumption needs to be very good, or the fit of the mean-variance relation needs to be bad, to justify using the gravity model as the lone source of extra information.

[1] http://www.cs.utexas.edu/users/yzhang/

3.10 Summary and Conclusions

In this chapter we studied the problem of traffic matrix estimation. Our work includes a theoretical statistical approach deriving performance bounds for estimators, a novel computationally light estimation method and a comparative evaluation of the effect that the nature of the additional information used has on estimation accuracy.

We derived an analytical expression for the Fisher information matrix in the traffic matrix estimation framework. This result was used to yield the Cramér-Rao lower bound for the variance of an estimator in the situation where we assume a functional mean-variance relationship for origin-destination flows in the network. We demonstrated why this result is extremely useful in various ways. We can obtain variances, and thus confidence intervals for the maximum likelihood estimate directly from the Cramér-Rao lower bounds. This means that we can identify the OD pairs whose estimates have large uncertainties. If the estimated traffic matrix is used, for instance, in load balancing, it should prove beneficial to know for which OD pairs the estimate might not be accurate. The CRLB can be used also in evaluation of estimation techniques, as we can compare the variance of the evaluated estimator to the lower bound to see how effective it is. The drawback of the bounds is that they assume that all assumptions of the maximum likelihood method hold. As we saw in the previous chapter, these assumptions are not 100% accurate. It is unclear how such inaccuracy affects the usefulness of the result. Also the bounds apply only to the second moment estimator framework, so comparison of several methods is not possible using the Cramér-Rao bounds.

We presented ways to obtain an estimate for a traffic matrix by explicit calculations utilizing the link count covariance matrix. We illustrated how one can obtain the OD-pair traffic variance estimates from an empirical link count covariance matrix, and developed computationally lightweight methods, the projection method and the constrained minimization method, to obtain an estimate for the traffic matrix based on the link count covariance matrix in a way that would still be consistent with the link counts. We evaluated the accuracy of the methods in a simulation study by comparing them against the maximum likelihood solution by Cao et al. The mean relative error for the quick method with projection was about fifty percent higher than it was for the maximum likelihood estimate. This is a reasonably good result, considering that the quick method is very simple to calculate, while the MLE is computationally too heavy for larger networks.

We proposed a novel estimation method, which combines the two competing sources of extra information: the gravity model and link covariances. Since typically both of them are relevant at least to some extent, as shown by the study of the three real data sets, it usually makes sense to use both. We showed that in many cases the combined method is the most accurate estimation technique.

Finally, we studied the sensitivity of the two most common methods to their underlying crucial assumptions. We found that the gravity method's accuracy deteriorates more quickly as a function of the error in gravity model fit than the second moment method does as a function of the fit of the mean-variance relation. However, when the assumptions hold, the gravity method is significantly more accurate. Also it needs only small sample sizes to achieve good accuracy, while the second moment method needs to estimate the sample covariance matrix, and thus is not very accurate with smaller samples. It would appear that, based on our study, the gravity-based methods are superior to the second moment methods with most realistic sample sizes. As for the combined method, we showed that it is typically rather accurate. The fit of the gravity model assumption needs to be very good, or the fit of the mean-variance relation needs to be bad, to justify using the gravity model as the lone source of extra information. If, for example, a large sample of, say, 500 measurements is available, the gravity model fit would have to be well over 0.90 to enable the gravity-model-based methods to outperform the combined method.

4 ROBUST LOAD BALANCING

4.1 Introduction

Load balancing is a common traffic engineering task in which the traffic in the network is routed in a way that optimizes some performance criterion. In this paper, we use the minimization of the maximal link utilization in the network as the target. This criterion leads to network usage that minimizes the relative congestion on the links throughout the network.

Traffic is moved from heavily congested links to other routes and links in order to ease the congestion in that part of the network. To do this, the traffic matrix is typically assumed to be known. However, as discussed in Chapter 3, traffic matrices are not generally readily obtainable. Estimated traffic matrices obviously are not entirely accurate, but come with some estimation error. Thus, the actual traffic load on a given link might significantly differ from what is expected on the basis of the estimated traffic matrix.

On the other hand, a method which does not require knowledge of the traffic matrix, or takes into consideration the estimation errors in the traffic matrix estimate, would be more robust. Robust approaches in general aim to achieve a situation where performance is good regardless of parameter values. For instance, if a system is dependent on a stochastic variable, we can optimize the system with all possible variable values in mind, not just the current one, or the most probable one.

In this case, we take into account the fact that the traffic matrix estimate is not accurate but contains error. The basic idea of the robust method is to balance the load in a way that does not optimize the network utilization for a single traffic matrix, but for a large polytope of matrices, which is selected so that it surely includes also the real traffic matrix.

In this chapter, we propose robust load balancing approaches for both IP and wireless multihop network frameworks.

4.2 Load Balancing

In this section we formulate the load balancing problem as well as the robust load balancing problem and define the objective function and variables.

The traditional load balancing approach relies on a known traffic matrix. We denote this by $\hat{\boldsymbol{x}}$, denoting that it is typically an estimate of the traffic matrix and not completely accurate. Let $\boldsymbol{L}$ denote the node-link incidence matrix with element $L_{nl} = +1$ and $L_{n'l} = -1$ if (directed) link l leads from node n to node n', and 0 otherwise, while $\boldsymbol{R}$ denotes the node - OD pair incidence matrix with element $R_{n_o,k} = +1$ and $R_{n_d,k} = -1$ if OD pair k enters the network at node n_o and exits at node n_d, and 0 otherwise.

Our performance metric u refers to the relative utilization of the most heavily congested link,

$$u = \max_l \frac{(\boldsymbol{Ax})_l}{C_l}. \tag{4.1}$$

Now the load balancing problem can be formulated as an LP problem

$$\min_{\boldsymbol{A} \geq \boldsymbol{0}} u \tag{4.2}$$

such that

$$u\boldsymbol{C} \geq \boldsymbol{A}\hat{\boldsymbol{x}}, \tag{4.3}$$

$$\boldsymbol{LA} = \boldsymbol{R}, \tag{4.4}$$

which yields the routing matrix that minimizes the maximum link utilization. The solution $\boldsymbol{A}(\hat{\boldsymbol{x}})$ is a function of the assumed traffic matrix $\hat{\boldsymbol{x}}$. This solution is optimal only with regard to the maximum link load, and usually is not unique. A second optimization is needed to ensure that traffic is optimally balanced in less loaded links also, not just in the bottleneck link.

4.3 Review of related work

Traffic engineering and load balancing have received lot of attention in the literature. For example MPLS networks give an excellent framework for load balancing applications [AMA+99] but also modifications to OPSF routing support traffic engineering [KKY03].

For instance Fortz et al. [FT00, FT02, FRT02] study extensively the approach to change link weights in order to balance the load in the network. Wang et al. [YW01] show that optimal routing can always be achieved by shortest path routing with some positive link weights if arbitrary split among the shortest paths is allowed. Fortz and Thorup extend their work in [FT03] to achieve robustness over link failures.

The first step in robust approach over demand uncertainty is the hose model proposed by Duffield et al. [DGG+99] and further developed by Erlebach and Ruegg [ER04]. In the hose model each endpoint of a Virtual Private Network (VPN) specifies bounds for its traffic demand. The provisioning is done so that there is sufficient bandwidth for any traffic matrix that is consistent with these specified bounds.

Ben-Ameur and Kerivin [BAK05] generalize the hose model and introduce the concept of routing a polytope, where the bounds for the possible traffic matrices are bounded on each link, not only the endpoints of the OD pairs. Johansson [Joh05] proposes the use of this concept for load balancing in the network without a traffic matrix estimate. While traditional load balancing relies on a single traffic matrix, the Robust method uses a polytope that contains all possible traffic matrices. The routing is then selected so that congested links are avoided for all traffic matrices of the polytope.

Applegate and Cohen [AC06] study the performance of robust routing for different size uncertainty sets. They construct uncertainty sets which are centered around the real demand and are sized by using multiples of the traffic matrix elements.

Roughan et al. [RTZ03] study the use of estimated traffic matrices in traffic engineering. They conclude that the usefulness of an estimator depends not only on the accuracy of the estimator, but also on the method that uses the estimated matrix in traffic engineering. A combination of tomogravity estimates and OSPF weight optimzation was found to work particularly well together.

Optimization of resource usage in wireless networks has also gained interest recently. Björklund et al. [BVY03] present mathematical programming approach for optimizing resource usage with STDMA by finding optimal time slot allocation.

Johansson and Xiao [JX06] propose a cross-layer optimization approach to find optimal routing, transmission schedule, power allocation for maximal throughput. Wu et al. [WCZ+05] also propose a cross-layer optimization approach. Their goal is to minimize congestion for multicast sessions. Jain et al. [JPPQ03] propose an interference aware routing protocol to achieve greater throughput. Within the proposed framework they are able to calculate lower and upper bounds for the optimal throughput.

Criteria other than throughput have also been addressed. Chen at al. [CWL05] aim for delay minimizations. Lassila et al. [LPV06] propose a cross-layer dimensioning approach for the wireless link capacities to fulfill flow level performance constraints. Susitaival [Sus07] proposes a load balancing method relying on joint optimization of routing and transmission schedule to achieve optimally balanced traffic load throughout the network.

4.4 Contribution of the Thesis

In Publication 7 we study the performance of robust load balancing methods against traffic matrix based load balancing methods. While traditional load balancing relies on a single traffic matrix and the Robust method on a polytope containing all possible traffic matrices, these are obviously extreme points. A middle ground between the two would be a polytope around the estimated traffic matrix, that is smaller than the set of all possible traffic matrices. Surely there are some cases so implausible that we do not need to consider them, and on the other hand it is unlikely that our traffic matrix estimate is so accurate that no error margins are needed. To obtain these bounded polytopes we propose a novel variant of the robust approach, the Bounded Robust Method. We analytically derive statistical standard error for the elements of the traffic matrix estimate and use different confidence intervals to obtain different size polytopes.

The differences in our approach compared to [AC06] is that we do not assume the actual traffic matrix and the estimated traffic matrix to coincide. We cannot construct the uncertainty set around the real demand, since it is unknown in reality. Only the estimate is available to us. Also, our error margins are statistical confidence intervals as opposed to multiples of the traffic amounts. As these confidence intervals are larger for origin-destination(OD) pairs that are difficult to estimate and smaller for those OD pairs that are easier to estimate, we avoid making the polytope unnecessarily large in directions where there is not that much uncertainty.

In Publication 8 we apply the robust approach to wireless mesh networks. We propose a robust approach which does not require knowledge of the traffic matrix, but only link count measurements and routing information. The idea behind the method is to balance the load for the worst case traffic matrix included in a larger polytope of matrices, not just for a single traffic matrix.

4.5 Robust Load Balancing with Estimated Traffic Matrices

In Publication 7 we introduce a novel approach for load balancing, the Bounded Robust Method, which combines the idea of the Robust method with the use of the traffic matrix estimate.

Robust Approach

In the Robust method, we assume that link count measurements and routing matrix is available, but do not use a traffic matrix estimate. Instead, we try to find a routing matrix such that the worst case performance is optimized over all feasible traffic matrices in the polytope

$$\mathcal{D} \quad = \quad \{\boldsymbol{x} \geq 0 : \boldsymbol{A}_0\boldsymbol{x} = \boldsymbol{y}_0\}. \tag{4.5}$$

This approach is proposed in [Joh05] and uses the algorithm introduced in [BAK05]. Our formulation is different from these in that we use the flow conservation constraints instead of making use of the arc-path formulation. The optimization problem in our framework is

$$\min_{\boldsymbol{A}\geq\boldsymbol{0}} u \tag{4.6}$$

such that

$$\begin{aligned} u\boldsymbol{C} \quad &\geq \quad \boldsymbol{A}\boldsymbol{x} \qquad \forall \boldsymbol{x} \in \mathcal{D} && (4.7)\\ \boldsymbol{L}\boldsymbol{A} \quad &= \quad \boldsymbol{R}. && (4.8) \end{aligned}$$

Again, a second step is needed to ensure that traffic is optimally balanced in less saturated links also.

This problem is difficult to solve because of the infinite number of constraints in (4.7). Therefore, a demand satellite approach [BAK05] has to be used. The problem is divided into two optimization problems. The first one is the *link load optimization problem*, where the set of infinite number of constraints $\mathcal{D}$ is substituted by a finite set of constraints D^*. These constraints are generated by the *constraint generation problem*. The link load optimization is now a simple LP problem, with the flow constraints just as in the traditional approach, and a set of finite constraints in the place of the link constraints.

Problem 3 (Link Load Optimization)

$$\min_{\boldsymbol{A}\geq\boldsymbol{0}} u \tag{4.9}$$

such that

$$
\begin{aligned}
u\,\boldsymbol{C}_l &\geq (\boldsymbol{A}\boldsymbol{x})_l \qquad \forall (l, \boldsymbol{x}) \in D^* && (4.10)\\
\boldsymbol{L}\boldsymbol{A} &= \boldsymbol{R}, && (4.11)
\end{aligned}
$$

and such that a secondary objective function is used to obtain optimal balancing throughout the network.

Problem 4 (Constraint Generation) For each link l solve with current values $\boldsymbol{A}^{(i)}$, $u^{(i)}$

$$\boldsymbol{x}^* = \arg\max_{x \in \mathcal{D}} (\boldsymbol{A}^{(i)}\boldsymbol{x})_l.$$

If

$$u^{(i)}\,\boldsymbol{C}_l < (\boldsymbol{A}^{(i)}\boldsymbol{x}^*)_l,$$

then

$$D^* \leftarrow D^* \cup (l, \boldsymbol{x}^*).$$

The set D^* is initially empty. Problem 4 is solved to obtain constraints. The iteration can be started using the initial routing $\boldsymbol{A}_0$. For each link we find the traffic matrix $\boldsymbol{x}^* \in \mathcal{D}$ that maximizes the traffic on that link. If the achieved link utilization is larger than the current value for u, the corresponding constraint

$$u\,\boldsymbol{C}_l \geq (\boldsymbol{A}\boldsymbol{x}^*)_l$$

is added to the set D^* to be used as a constraint in (4.10) in the next iteration of problem 3.

The Bounded Robust Method

In the constraint generation problem, the new constraints are obtained by finding the traffic matrix $\boldsymbol{x}^* \in \mathcal{D}$ that maximizes traffic on a particular link. This traffic matrix is typically one which has extreme values for OD pairs that use this particular link and zero values for most other OD pairs. These cases are in the set $\mathcal{D}$, but in reality, are very unlikely. If we are interested in a plausible maximum link load, we might like to add some more constraints to eliminate the cases which are practically impossible. This can be achieved by setting an upper bound on the value that an estimate for an OD pair may obtain.

In Publication 7 we propose the Bounded Robust algorithm that is based on the idea that we do utilize a traffic matrix estimate, but also take into consideration the error of the estimator. We calculate confidence intervals for the estimates using the Cramér-Rao lower bounds (CRLB) for the variance of an estimator, derived in Publication 4.

We redefine the set $\mathcal{D}$ of traffic matrices to take into consideration these confidence intervals. Note that we have no need to bound the lower values, as we are looking for the maximal link loads. Thus we define this new set as

$$\mathcal{D}' = \{\boldsymbol{x} : \boldsymbol{A}_0\boldsymbol{x} = \boldsymbol{y}_0\ ,\ 0 \leq \boldsymbol{x} \leq \hat{\boldsymbol{x}} + z \cdot \boldsymbol{S}_E\}, \tag{4.12}$$

and the constraint generation problem is changed accordingly. An iteration is performed as described for the robust method in section 4.5, with the new constraint generation problem now in place. The result of this algorithm gives the optimal routing matrix over all reasonably conceivable traffic matrices.

Simulation Results

Robust methods by definition have lower worst case link utilizations. But this provisioning for the worst case might have an adverse effect on the mean utilization obtained. We evaluate this tradeoff between average link utilization and robustness by a simulation study comparing the following approaches

1. The traditional load balancing with maximum likelihood traffic matrix estimate,
2. The Robust method,
3. The Bounded Robust method.

To evaluate the performance of each algorithm, we define three performance metrics.

1. The real case utilization of the most congested link.

   $$u_{\mathrm{rc}} = |\widehat{\boldsymbol{A}}\boldsymbol{x}_0/\boldsymbol{C}|, \tag{4.13}$$

   where x_0 is the real value of the traffic matrix during the measurements.

2. The worst case maximal utilization.

   $$u_{\mathrm{wc}} = \max_{\boldsymbol{x}\in\mathcal{D}} |\widehat{\boldsymbol{A}}\boldsymbol{x}/\boldsymbol{C}|. \tag{4.14}$$

   We consider the largest value for u with the routing matrix $\widehat{\boldsymbol{A}}$, changing the traffic matrix within the polytope specified in (4.5). This is the criterion the Robust methods aims to minimize.

3. The bounded worst case utilization is similar to the second one except that we exclude implausible traffic matrices from consideration by finding the maximal u such that the traffic matrix is within the bounded polytope given in (4.12).

   $$u_{\mathrm{bwc}} = \max_{\boldsymbol{x}\in\mathcal{D}'} |\widehat{\boldsymbol{A}}\boldsymbol{x}/\boldsymbol{C}|. \tag{4.15}$$

   This is the criterion the Bounded Robust methods aims to minimize.

All the results in the sequel are given relative to the optimal load balancing that would be obtained using the exact traffic matrix for the traditional load balancing method.

Figure 4.1 depicts the real case utilization and the worst case utilization obtained by the different methods. The utilization is lowest with the traditional method with $u_{rc} = 1.04$. For a small confidence interval with coefficient $z = 0.5$ it is still quite low, but then grows quickly when the

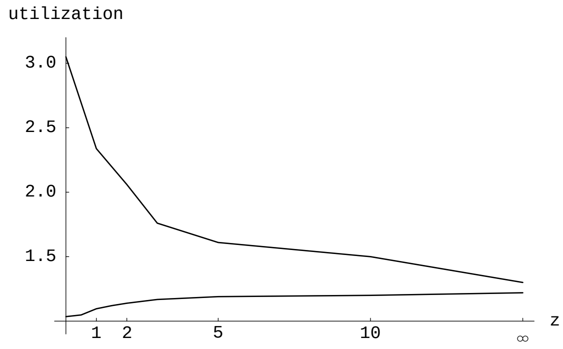

Figure 4.1: Real case performance (lower curve) and worst case performance (upper curve) of the Bounded Robust method as a function of the confidence interval coefficient

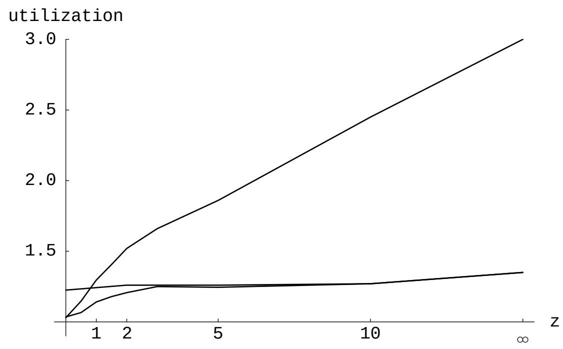

Figure 4.2: Bounded worst case performance of Traditional (the fast rising curve), Bounded Robust (lower curve) and Robust methods (middle curve) as a function of the confidence interval coefficient.

multiplying coefficient z approaches 3. After that the utilization grows only moderately from 1.17 to 1.22 which is the real case utilization for the Robust method.

Figure 4.2 shows the bounded worst case utilizations for the different methods. On the horizontal axis is the width of the confidence interval, over which the bounded worst case utilization u_{bwc} is calculated. So, $z = 0$ corresponds to using no confidence interval at all. In this case u_{bwc} is exactly the same as the real case utilizations, as we calculate the worst case over a zero size polytope, which is the single point of the estimate. Then, for instance at the point $z = 1.96$ the curves give the bounded worst case utilizations for the polytope of 95% confidence intervals. Finally, at $z = \infty$, u_{bwc} coincides with u_{wc} as the confidence interval is so large that we take the worst case over all possible traffic matrices. The highest curve is the traditional approach. The curve for Bounded Robust starts lower than the Robust, but they approximately coincide after $z = 3$.

Table 4.1: Values of utilization for different methods when $z = 1.96$.

	u_{rc}	u_{bwc}	u_{wc}
Traditional	1.04	1.52	3.05
Bounded Robust	1.13	1.20	2.06
Robust	1.22	1.26	1.31

For example, let us consider the 95% confidence interval for each element of the traffic matrix. We have three methods: Traditional method, Bounded Robust method with $z = 1.96$ and Robust method. We can read from Figure 4.2 the bounded worst case utilization for each. From Figure 4.1 we can read the real case and worst case utilizations using $z = 0$ for Traditional method, $z = 1.96$ for Bounded Robust and $z = \infty$ for Robust. These results are listed in Table 4.1.

The Bounded Robust method achieves better performance for all plausible traffic matrices compared to the Robust method. By eliminating the cases that are, though still in line with the measurements, extremely unlikely, it is possible to obtain clearly better performance for the most probable cases, as well as slightly better for the range on plausible cases.

4.6 Robust Load Balancing in Wireless Networks

Publication 8 studies load balancing in a wireless mesh network [AW05]. We present a robust approach that enables cross-layer optimization of both transmission schedule and routing without the knowledge of traffic matrix.

Modeling the Wireless Network

The MAC layer is modeled by Spatial TDMA [NK85], where the transmission resources are divided into time slots. The links that are not interfering each other can transmit in the same time slot. A set of links that are able to transmit simultaneously is called a transmission mode.

The interference model assumes that interference restricts the links that can transmit simultaneously, but does not affect capacity of those that are able to transmit. The communication range of node i is denoted by R_i and the distance between nodes i and j by d_{ij}. Now, if

$$d_{ij} \leq R_i$$

there is a link between nodes i and j. The interference range is denoted by R'_i. For transmission to be successful on link ij, there cannot be a node k transmitting such that

$$d_{kj} \leq R'_k.$$

Based on this criterion we can determine all possible transmission modes. That is, the sets of links that can transmit simultaneously. Finding all feasible modes requires extensive calculations, but if the topology is not large, the maximal transmission modes can be generated by the algorithm proposed in [Sus07].

Let $\boldsymbol{b}$ be a vector of length L denoting the nominal link bandwidths. The elements of the vector are the link bandwidths for each link, with the links indexed by l. There are M different transmission modes, which are represented in the $L \times M$ matrix $\boldsymbol{S}$. It has an element $S_{l,m} = b_l$ if link l transmits in mode m and $S_{l,m} = 0$ otherwise.

The M-vector $\boldsymbol{t}$, whose elements sum to unity, denotes the fraction of time spent in each transmission mode. The actual capacities of the links is then denoted by

$$\boldsymbol{C} = \boldsymbol{S}\boldsymbol{t}. \tag{4.16}$$

Further, we define vector $\boldsymbol{q}$ such that

$$q_m = u t_m, \tag{4.17}$$

where the elements of $\boldsymbol{q}$ sum to u, which represents the fraction of time spent in any transmission mode, with the rest of the time the network being idle.

As in the previous section, we denote the polytope of feasible traffic matrices by $\mathcal{D}$.

MAC Layer Optimization

Optimizing only the MAC layer transmission schedule based on the link counts is a straightforward approach to achieve load balancing in the network without the knowledge of the traffic matrix. In essence, the link capacities are changed to fit the traffic volumes on the links by changing the scheduling. We denote the current routing matrix, for instance the shortest path routing, by $\boldsymbol{A}_0$, and the corresponding link counts are denoted by $\boldsymbol{y}_0$. The optimization problem is

Problem 5 (MAC Layer Optimization)

$$\min_{\boldsymbol{q} \geq \boldsymbol{0}} \boldsymbol{e}^T \boldsymbol{q} \tag{4.18}$$

such that

$$\boldsymbol{S}\,\boldsymbol{q} \quad \geq \quad \boldsymbol{y}_0. \tag{4.19}$$

The result gives the optimal schedule for the fixed routing which yielded the link counts $\boldsymbol{y}_0$.

Optimizing the Layers Separately

One load balancing approach would be to first balance the load using the nominal link capacities, and then optimize the transmission schedule. However, as we assume that the traffic matrix is not available, we cannot use the traditional approach for the load balancing. Instead, we have to use the robust approach. Another major drawback is that we would then need to make new measurements to obtain the link counts associated with the new routing matrix in order to optimize the transmission schedule.

Cross-Layer Robust Load Balancing

To optimize the layers at the same time we have to modify the approach of section 4.5 to wireless framework by incorporating the transmission schedule variables $\boldsymbol{q}$ to the optimization problem. The minimization problem is now

Problem 4 (Cross-layer Robust Load Balancing)

$$\min_{\boldsymbol{A}\geq\boldsymbol{0}\,,\,\boldsymbol{q}\geq\boldsymbol{0}} \boldsymbol{e}^T\boldsymbol{q} \tag{4.20}$$

such that

$$\boldsymbol{S}\,\boldsymbol{q} \geq \boldsymbol{A}\boldsymbol{x} \qquad \forall \boldsymbol{x}\in\mathcal{D}, \tag{4.21}$$

$$\boldsymbol{L}\boldsymbol{A} = \boldsymbol{R}. \tag{4.22}$$

Again, we use the iterative approach described above.

Simulation Results

To validate the above approach we performed a simulation study with synthetic traffic, using a 12-node topology.

Table 4.2 compares the results of the robust methods, that do not use a traffic matrix. The results are normalized such that 1.00 is the theoretical optimal, which is the value obtained by the cross-layer optimization using the accurate traffic matrix. The initial utilization is given by using the shortest path routing and a transmission schedule that assigns equal time for each transmission mode.

Table 4.2: Values of u for different methods, with 1.00 being the optimal solution

Initial	3.17
MAC layer only (shortest path routing)	1.19
Two-layer robust approach	1.18
Cross-layer robust approach	1.10

It can be seen that the cross-layer approach performs clearly best of the robust methods. Although it has to optimize the routing for all traffic matrices in the polytope, there is still enough information about the traffic matrix so that it is beneficial to optimize routing also instead of just the transmission schedule. The two-layer approach is only marginally more accurate than the MAC layer only optimization and thus probably not useful.

If a traffic matrix estimate is available, the accuracy of an optimization based on this estimate would obviously depend on the accuracy of the estimate. Therefore we evaluated the performance of this method using different estimates. Figure 4.3 shows how the performance of this approach

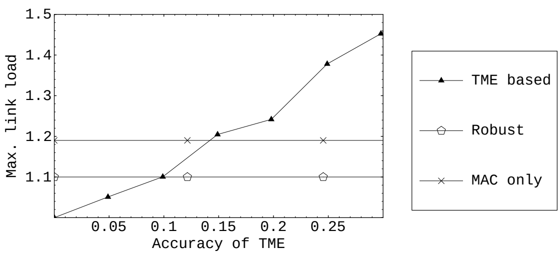

Figure 4.3: Performance of traffic matrix optimization approaches.

deteriorates compared to the cross-layer Robust and MAC layer only approaches as the accuracy of the estimate gets worse. It can be seen that the approach using estimates quickly becomes worse than the robust approaches. With an estimation error larger than 10% it is already worse than the cross-layer robust approach. This size of error is considered typically rather small in traffic matrix estimation. In fact, in IP networks this is often considered the target value which estimation methods strive for. Only third generation methods that use extensive Netflow measurements achieve typically errors from 5% to 15%, while traditional methods relying on SNMP measurements cannot usually do better than 20% [SLT+05].

4.7 Summary and Conclusions

In this chapter, we address new approaches for the network load balancing task. Instead of relying on an estimated and inaccurate traffic matrix, we study methods that use the robust approach, where the optimization is made so that the performance is satisfactory in any situation. This is achieved by optimizing the maximal link load over all possible, or, alternatively, all plausible, traffic matrices.

In fixed networks, this robust approach, which does not need any traffic matrix at all, was found to perform sufficiently well. The strength of the approach is that the worst case link utilizations is only 31% higher than optimal, while the worst case for the Traditional method is over three times the optimal utilization.

In Publication 7, we proposed a novel extension for the robust approach: the Bounded Robust method, which requires a traffic matrix estimate, but takes the uncertainty in the estimation into account. This approach decreases the size of the polytope of considered traffic matrices by introducing confidence intervals for the estimator, thus eliminating the need to include the next-to-impossible extreme cases in the provisioning. The method uses maximum likelihood estimates for the traffic matrix and makes use of the Cramér-Rao lower bounds to obtain the confidence intervals. In some simulation scenarios the Bounded Robust method was shown not only to add the robustness that the traditional approaches lack, but actually outperform them with regard to the real case maximal link utilization. In all situations the method is close to the traditional method and outperforms the Robust method in real case link utilization, while still maintaining the robust performance over all plausible traffic matrices.

While the results are encouraging, it may be impossible to implement this kind of approach in current networks, where the routing possibilities and use of excessive multipath routing are not available. However, this approach does give an idea of the theoretical possibilities of such a routing scheme.

In Publication 8, we adapted the robust load balancing approach to wireless networks. We proposed a cross-layer approach, where the transmission schedule and routing are optimized simultaneously.

We studied the performance of the methods by a simulation study and found that the cross-layer robust method was the best of the studied methods. It yielded maximal link utilization of only 10% worse than optimal and outperformed the other approaches. It was shown that the traditional load balancing method using an estimated traffic matrix performed worse than the cross-layer robust method unless it had an exceptionally accurate traffic matrix estimate.

In the light of these results, it is highly questionable whether it is worth the trouble to try to estimate the traffic matrix in a wireless environment, since in real networks there is no way of knowing the actual accuracy of the given estimate, and most likely the robust approach outperforms the method using the estimate.

However, we must note that this method is a straightforward generalization of the robust method for fixed networks. It depends on the fact that link measurements and routing tables are readily available. While this is true in fixed IP networks, in wireless multihop networks, the nodes might disappear or new nodes may appear, so it is not clear whether such information that is in line with the current state of the network is always available.

5 SUMMARY

This thesis studied the problem of traffic matrix estimation. The scope covered Internet traffic characterization, different techniques for traffic matrix estimation, performance evaluation of these estimation methods with simulation studies, as well as analytical calculations and load balancing applications of the estimation methods.

The traffic characterization part of the thesis was based on the study of the Finnish University and Research Network traffic measurements. The traffic on a backbone link between Helsinki and Espoo was captured and analyzed. We analyzed the link aggregate traffic and also divided the trace into origin-destination pair traffic for several different aggregation levels. Each of these data sets was analyzed under different time-scale granularities. This enabled a thorough characterization of OD-pair traffic qualities and of how the traffic characteristics and validity of assumptions about the nature of the traffic behave when the aggregation level changes both temporally and spatially. We were able to confirm that the Gaussian distribution is a relatively good fit for the stochastic fluctuation of the traffic, noting that, the larger the traffic flow, the better the fit of the distribution. We also confirmed that the assumption about a functional relation between the mean and the variance of the traffic is at least reasonable, while not perfectly accurate. These results are important in the context of traffic matrix estimation, as the second moment methods would be worthless without the assumption about a mean-variance relation and the maximum likelihood method needs an accurate traffic distribution model with which to construct the likelihood equation.

The derivation of the Cramér-Rao lower bounds was one of the key results of the second part of the thesis. The bounds give the lowest possible variance for the non-biased second moment estimator. This makes it possible to calculate the best accuracy that can be achieved by the estimator in a given situation. The maximum likelihood estimator is efficient and its variance coincides with the bound. The Cramér-Rao lower bounds are an important theoretical concept, but we showed that they also have many practical uses, such as finding the optimal places for direct measurements. We also utilized the bounds in a load-balancing problem later in the thesis.

Concerning the estimation methods, the key observation that most of the methods can be classified into two classes based on the source of extra information used emerged. One class uses the gravity model, which assumes that traffic between two given nodes is proportional to the total traffic originating from the source node and terminating at the destination node. These assumptions restrict the possible traffic matrix values enough to make the system identifiable. The other class uses the variance of the traffic as the extra information by utilizing a functional mean-variance relation to make use of the variances in the estimation of the mean. The variances are solvable from the link sample covariances. Thus there are basically two approaches: the gravity model estimators and the second moment estimators. Both approaches also obviously take the measured link loads into account. A key difference is that the gravity model methods only need a single measurement snapshot of the network to yield an estimate. The second moment methods, on the other hand, need a time series of locally stationary measurement samples, which has to be large enough to yield reliably accurate sample covariances.

The differences between methods within the same class are mostly in the mathematical implementation. The exception to this are the approaches called *third generation* methods, which use occasional direct OD-pair measurements as the source of the extra information and thus cannot be classified into the classes of the traditional methods.

The thesis also proposed two novel estimation methods. Both of these are computationally fast but yield reasonably accurate results. The idea behind the quick method was to make use of the assumptions in the maximum likelihood estimation framework while getting rid of the exhaustive calculations that are too time consuming for large networks. This was achieved by sacrificing some of the accuracy of the maximum likelihood method. We used a two-stage method where the first part obtains a prior estimate using the link covariances and mean variance relation, while the second part projects this estimate onto the plane where the link count equation holds.

The combined method is the first step in combining the two classes of extra information that had been used separately thus far. The method itself is as simple as the quick method, but the idea of incorporating relevant information from both sources might lead to something more substantial.

The final part draws from a fact that was evident throughout the thesis: the traffic matrices we have available for traffic engineering are typically just estimates. Estimation always includes error, and thus traffic engineering tasks performed with estimated traffic matrices should take this error into account. We therefore studied robust approaches, where the performance is optimized for a range of possible traffic matrices, not just a single estimate. First, we studied the robust approach for fixed networks. Robust approaches in the literature guaranteed robust performance over all possible traffic matrices. We introduced a new approach where the robustness is only over a polytope of plausible traffic matrices. We used the Cramér-Rao bounds to construct this polytope, binding the second and third parts of the thesis even more closely together. Secondly, we applied the robust approach to wireless networks and proposed a robust cross-layer optimization approach for load balancing.

The thesis spanned the traffic engineering framework from measurements to load balancing. Our emphasis was on the traffic matrix estimation. We can conclude that the estimation problem is a difficult one because of its highly underconstrained nature. The assumptions that provide the extra information were validated to be reasonable but not totally accurate. This introduces inaccuracy and bias to the estimators. Adding to this the fact that any estimator of a stochastic variable has some minimum variance and thus comes with an estimation error, it is safe to say that an accurate traffic matrix often assumed in traffic engineering is not obtainable in reality. Having said this, we must note that the estimates are by no means useless. The estimators proposed in this thesis provide, without heavy calculations, a fairly accurate estimate of the traffic matrix. Some of the methods in the literature that have a longer running time yield even more-accurate estimates. This can very well be used in traffic engineering tasks. However, it would seem sensible to recommend that the inaccuracy of the estimator be taken into account, or perform some sort of sensitivity analysis when designing applications that use an estimated traffic matrix.

6 AUTHOR'S CONTRIBUTION

Publication 1
This paper is joint work by the authors. Markus Peuhkuri is responsible for performing the measurements and the preliminary classification of the data. Dr. Samuli Aalto wrote the introduction and preliminary sections. The present author and Riikka Susitaival analyzed the data and wrote the rest of the paper.

Publication 2
This paper is joint work by the authors. Markus Peuhkuri is responsible for performing the measurements and the preliminary classification of the data. Dr. Samuli Aalto wrote the introduction and preliminary sections. The present author and Riikka Susitaival analyzed the data and wrote the rest of the paper. The present author is responsible for analyzing the mean-variance relation.

Publication 3
This paper is joint work by the authors. Markus Peuhkuri is responsible for performing the measurements and the preliminary classification of the data. The present author and Riikka Susitaival analyzed the data and wrote the paper. The present author is responsible for analyzing the mean-variance relation and studying the moving Gaussian model.

Publication 4
This paper is joint work by the authors. The idea is by Prof. Sandrine Vaton. The calculations were made by Paola Bermolen. The present author is responsible for the applications section and writing of the paper.

Publication 5
The idea for the method proposed in the paper is by the present author. The computational techniques were developed as joint work by the authors. The present author made the numerical evaluations and wrote the paper.

Publication 6
This paper is an independent work by the author.

Publication 7
This paper is an independent work by the author.

Publication 8
This paper is an independent work by the author.

[AC06] D. Applegate and E. Cohen. Making intra-domain routing robust to changing and uncertain traffic demands: Understanding fundamental tradeoffs. *IEEE/ACM Transactions on Networking, Vol.* 14*, No.* 6, 2006.

[AMA+99] D. Awduche, J. Malcolm, J. Agogbua, M. O'Dell, and J. McManus. Requirements for traffic engineering over mpls, September 1999. IETF RFC 2702.

[AW05] I. Akyildiz and X. Wang. A survey on wireless mesh networks. *IEEE Radio Communications*, september 2005.

[BAK05] W. Ben-Ameur and H. Kerivin. Routing of uncertain traffic demands. *Optimization and Engineering*, 6(3), September 2005.

[BDJT01] S. Bhattacharyya, C. Diot, J. Jetcheva, and Nina Taft. Pop-level and access-link-level traffic dynamics in a tier-1 pop. In *ACM Internet Measurement Workshop (IMW)*, 2001.

[BVY03] P. Björklund, P. Värbrand, and D. Yuan. Resource optimization of spatial tdma in ad hoc radio networks: A column generation approach. In *IEEE Infocom*, 2003.

[CB96] M. Crovella and A. Bestavros. Self-similarity in world wide web traffic: Evidence and possible causes. In *SIGMETRICS'96*, Philadelphia, United States, May 1996.

[CCLS01] J. Cao, W. Cleveland, D. Lin, and D. Sun. On the nonstationarity of internet traffic. In *SIGMETRICS/Performance*, pages 102–112, 2001.

[CDWY00] J. Cao, D. Davis, S. V. Wiel, and B. Yu. Time-varying network tomography. *Journal of the American Statistical Association*, 95:1063–1075, 2000.

[CFSD90] J.D. Case, M. Fedor, M.L. Schoffstall, and C. Davin. Simple network management protocol (snmp). In *RFC1157, STD0015*, 1990.

[CFT+02] P. Carlsson, M. Fiedler, K. Tutschku, S. Chevul, and A. Nilsson. Obtaining reliable bit rate measurements in snmp-managed networks. In *ITC Specialist Seminar*, Wurzburg, Germany, 2002.

[Cis] Cisco netflow, www.cisco.com/warp/public/732/tech/netflow.

[CL03] A.E. Conway and M. Li. Estimating the traffic distribution matrix of a packet-switched network using aggregate traffic measurements. *Telecommunication Systems*, 24(1), 2003.

[CTB98] M. Crovella, M. Taqqu, and A. Bestavros. Heavy-tailed probability distributions in the world wide web, 1998.

[CWL05] C. Chen, W. Wu, and Z. Li. Multipath routing modelling in ad hoc networks. In *ICC*, Seoul, Korea, 2005.

[CWYZ01] J. Cao, S.V. Wiel, B. Yu, and Z. Zhu. A scalable method for estimating network traffic matrices. In *Bell Labs Technical Report*, 2001.

[DGG+99] N. Duffield, P. Goyal, A. Greenberg, P. Mishra, K. K. Ramakrishnan, and J. van der Merive. A flexible model for resource management in virtual private networks. In *SIGCOMM*, pages 95–108, 1999.

[EMH05] S. Eum, J. Murphy, and R.J. Harris. A fast accurate lp approach for traffic matrix estimation. In *ITC 19*, Beijing, China, 2005.

[ER04] T. Erlebach and M. Ruegg. Optimal bandwidth reservation in hose-model vpns with multi-path routing. In *in proceedings on Infocom 2004*, 2004.

[FGL+01] A. Feldmann, A. Greenberg, C. Lund, N. Reingold, J. Rexford, and F. True. Deriving traffic demands for operational IP netowrks: Methodology and experience. *IEEE/ACM Transactions on Networking*, 9(3), 2001.

[FRT02] B. Fortz, J. Rexford, and M. Thorup. Traffic engineering with traditional ip routing protocols. *IEEE Communications Magazine*, October 2002.

[FT00] B. Fortz and M. Thorup. Internet traffic engineering by optimizing OSPF weights. In *INFOCOM (2)*, pages 519–528, 2000.

[FT02] B. Fortz and M. Thorup. Optimizing ospf/is-is weights in a changing world. *IEEE Journal on Selected Areas in Communications*, 20, 2002.

[FT03] B. Fortz and M. Thorup. Robust optimization of ospf/is-is weights. In *INOC 2003*, 2003.

[GJT04] A. Gunnar, M. Johansson, and T. Telkamp. Traffic matrix estimation on a large ip backbone – a comparison on real data. In *IMC'04*, Taormina, Italy, 2004.

[Gol00] O. Goldschmidt. Isp backbone traffic inference methods to support traffic engineering. In *Internet Statistics and Metrics Analysis Workshop*, San Diego, USA, 2000.

[Joh05] M. Johansson. Data driven traffic engineering. Presentation in LSNI Workshop, 2005.

[JPPQ03] K. Jain, J. Padhye, V. Padmanabhan, and L. Qiu. Impact of interference on multi-hop wireless network performance. In *Internationall Conference on Mobile Computing and Networking*, San Diego, United States, September 2003.

[JX06] M. Johansson and L. Xiao. Cross-layer optimization of wireless networks using nonlinear column generation. *IEEE Transactions on Wireless Communications*, 5(2), february 2006.

[KKY03] D. Katz, K. Kompella, and D. Yeung. Traffic engineering (te) extensions to ospf version 2, September 2003. IETF RFC 3630.

[KMFB04] T. Karagiannis, M. Molle, M. Faloutsos, and A. Broido. A nonstationary poisson view of internet traffic. In *IEEE Infocom*, Hong Kong, 2004.

[KN02] J. Kilpi and I. Norros. Testing the gaussian approximation of aggregate traffic. In *ACM SIGCOMM Internet Measurement Workshop*, Marseille, France, 2002.

[KW95] J. Kowalski and B. Warfield. Modeling traffic demand between nodes in a telecommunications network. In *ATNAC 95*, 1995.

[LaJW97] D. Lam and D. Cox adn J. Widom. Teletraffic modelling for personal communications services. *IEEE Communications Magazine: Special Issues on Teletraffi Modellinf Engineering and Management in Wireless and Broadband Networks*, 35, Ferbruary 1997.

[LPC$^+$04] A. Lakhina, K. Papagiannaki, M. Crovella, C. Diot, E.D. Kolaczyk, and N. Taft. Structural analysis of network traffic flows. In *SIGMETRICS/Performance'04*, New York, USA, 2004.

[LPV06] P. Lassila, A. Penttinen, and J. Virtamo. Dimensioning of wireless mesh networks. In *ACM PE-WASUN'06*, Malaga, Spain, 2006.

[LTWW94] W. Leland, M. Taqqu, W. Willinger, and D. Wilson. On the self-similar nature of ethernet traffic (extended version). *IEEE/ACM Transactions on Networking*, 2(1), 1994.

[LTY06] G. Liang, N. Taft, and B. Yu. A fast lightweight approach to orgin-destination ip traffic estimation using partial measurements. *Special joint issue of IEEE Transactions on Networking and IEEE Transactions on Information Theory*, June 2006.

[Lue69] D. Luenberger. *Optimization by Vector Space Methods*. John Wiley & Sons, New York NY, 1969.

[LW91] W. Leland and D. Wilson. High time-resolution measurement and analysis of LAN traffic: Implications for LAN interconnection. In *INFOCOM (3)*, pages 1360–1366, 1991.

[LY03] G. Liang and B. Yu. Pseudo likelihood estimation in network tomography. In *IEEE Infocom*, 2003.

[MD00] S. Molnar and T. Dang. Pitfalls in long range dependence testing and estimation, 2000.

[MK97] G. J. McLachlan and T. Krishnan. *The EM Algorithm and Extensions*. John Wiley & Sons, New York NY, 1997.

[ML00] R. Morris and D. Lin. Variance of aggregated web traffic. In *IEEE Infocom*, Tel Aviv, Israel, 2000.

[MST+03] A. Medina, K. Salamatian, N. Taft, I. Matta, Y. Tsang, and C. Diot. On the convergence of statistical techniques for inferring network traffic demands. In *Technical Report*, 2003.

[MST+04] A. Medina, K. Salamatian, N. Taft, I. Matta, and C. Diot. A two-step statistical approach for inferring network traffic demands. In *Technical Report*, 2004.

[MTS+02] A. Medina, N. Taft, K. Salamatian, S. Bhattacharyya, and C. Diot. Traffic matrix estimation: Existing techniques and new solutions. In *SIGCOMM'02*, Pittsburg, USA, 2002.

[NCTD04] A. Nucci, R. Cruz, N. Taft, and C. Diot. Design of igp link weight changes for estimation of traffic matrices. In *IEEE Infocom*, Hong Kong, China, 2004.

[NK85] R. Nelson and L. Kleinrock. Spatial-tdma: A collision-free multihop channel access control. *IEEE Transactions on Communications*, 33, 1985.

[Nor94] I. Norros. A storage model with self-similar input. *Queueing Systems*, 16, 1994.

[Nor95] I. Norros. The management of large flows of connectionless traffic on the basis of self-similar modeling, 1995.

[Par97] K. Park. On the effect and control of self-similar network traffic: A simulation perspective. In *Winter Simulation Conference*, pages 989–996, 1997.

[Pax97] V. Paxson. *Measurement and Analysis of End-to-End Internet Dynamics*. PhD thesis, University of California, Berkeley, 1997.

[Peu01] M. Peuhkuri. A method to compress and anonymize packet traces. In *IMW'01*, San Fransisco, USA, 2001.

[PF95] V. Paxson and S. Floyd. Wide area traffic: the failure of poisson modeling. *IEEE/ACM Transactions on Networking*, 3(3):226–244, 1995.

[PKC96] K. Park, G. Kim, and M. Crovella. On the relationship between file sizes, transport protocols, and self-similar network traffic. In *Proceedings of IEEE International Conference on Network Protocols*, pages 171–180, October 1996.

[PTL04] K. Papagiannaki, N. Taft, and A. Lakhina. A distributed approach to measure traffic matrices. In *ACM Internet Measurement Conference*, Taormina, Italy, 2004.

[Pyh63] P. Pyhnen. A tentative model for the volume of trade between countries. *Weltwirtschaftliches Archive*, 90, 1963.

[RC99] C. Robert and G. Casella. *Monte Carlo Statistical Methods*. Springer, 1999.

[RSCA06] M. Rahman, S. Saha, U. Chengan, and A.S. Alfa. Ip traffic matrix estimation methods: Comparisons and improvements. In *ICC 2006*, Istanbul, Turkey, 2006.

[RTZ03] M. Roughan, M. Thorup, and Y. Zhang. Traffic engineering with estimated traffic matrices. In *USENIX/ACM Internet Measurement Conference*, Miami, United States, October 2003.

[RVC01] E. Rosen, A. Viswanathan, and R. Callon. Multiprotocol label switching architecture. In *IETF RFC 3031*, January 2001.

[Sha98] J. Shao. *Mathematical Statistics*. Springer, New York, 1998.

[SLT+05] A. Soule, A. Lakhina, N. Taft, K. Papagiannaki, K. Salamatian, A. Nucci, M. Crovella, and C. Diot. Traffic matrices: Balancing measurements, inference and modeling. In *SIGMETRICS'05*, Banff, Canada, 2005.

[SNC+04] A. Soule, A. Nucci, R. Cruz, E. Leonardi, and N. Taft. How to identify and estimate the largest traffic matrix elements in a dynamic environment. In *SIGMETRICS/Performance'04*, New York, USA, 2004.

[SSNT05] A. Soule, K. Salamatian, A. Nucci, and N. Taft. Traffic matrix tracking using kalman filtering. In *ACM SIGMETRICS Performance Evaluation Review (PER)*, volume 33, 2005.

[Sus07] R. Susitaival. Load balancing by joint optimization of routing and scheduling in wireless mesh networks. In *ITC-20*, Ottawa, Canada, 2007.

[Tin62] J. Tinbergen. Shaping the world economy: Suggestions for an international economy policy. In *The Twnetieth Century Fund*, 1962.

[TW98] C. Tebaldi and M. West. Bayesian inference on network traffic using link count data. *Journal of the American Statistical Association*, pages 557–576, 1998.

[Var96] Y. Vardi. Network tomography: estimating source-destination traffic intensities from link data. *Journal of the American Statistical Association*, pages 365–377, 1996.

[VBG05] S. Vaton, J.S. Bedo, and A. Gravey. Advanced methods for the estimation of the origin destination traffic matrix. In *Revue du 25*ème *anniversaire du GERAD*, 2005.

[vdMMP06] R. van de Meent, M. Mandjes, and A. Pras. Gaussian traffic everywhere? In *ICC 2006*, Istanbul, Turkey, 2006.

[VG03] S. Vaton and A. Gravey. Network tomography: an iterative bayesian analysis. In *ITC18*, Berlin, Germany, 2003.

[WCZ+05] Y. Wu, P. Chou, Q. Zhang, K. Jain, W. Zhu, and S. Kung. Network plannign in wireless ad hoc networks: A cross-layer approach. *IEEE Journal on Selected Areas in Communications*, 23(1), 2005.

[YW01] L. Zhang Y. Wang, Z. Wang. Internet traffic engineering without full mesh overlaying. In *Infocom*, 2001.

[ZRDG03] Y. Zhang, M. Roughan, N. Duffield, and A. Greenberg. Fast accurate computation of large-scale ip traffic matrices from link loads. In *ACM Sigmetrics*, 2003.

[ZRLD03] Y. Zhang, M. Roughan, C. Lund, and D. Donoho. An information-theoretic approach to traffic matrix estimation. In *ACM SIGCOMM*, 2003.