Mälardalen University Press Dissertations
No. 42

MODELING, IMPLEMENTATION AND EVALUATION OF IP NETWORK BANDWIDTH MEASUREMENT METHODS

Andreas Johnsson

2007

MÄLARDALEN UNIVERSITY

Department of Computer Science and Electronics
Mälardalen University

ISSN 1651-4238
ISBN 978-91-85485-41-3
Printed by Arkitektkopia, Västerås, Sweden

Mälardalen University Press Dissertations
No. 42

MODELING, IMPLEMENTATION AND EVALUATION OF IP NETWORK BANDWIDTH MEASUREMENT METHODS

Andreas Johnsson

Akademisk avhandling

som för avläggande av Teknologie doktorsexamen i Datavetenskap vid Institutionen för datavetenskap och elektronik kommer att offentligen försvaras fredagen, 27:e april, 2007, 13.15 i Gamma, U, Rosenhill.

Fakultetsopponent: professor Gábor Vattay, Collegium Budapest.

Institutionen för datavetenskap och elektronik
Mälardalen University, Box 883, SE-72123 Västerås

Abstract

Internet has gained much popularity among the public since the mid 1990's and is now an integrated part of our society. A large range of high-speedbroadband providers and the development of new and more efficient Internet applications increase the possibilities to watch movies and live TV, use IP-telephony and share files over the Internet. Such applications demand high data transmission rates, which in turn consume network bandwidth. Since several users must share the common bandwidth capacity on the Internet, there will be locations in the network where the demand is higher than the capacity. This causes network congestion, which has negative impact on both the data transmission rate and transmission quality.

This thesis is about methods for measuring the available bandwidth of a network path between two computers. The available bandwidth can be interpreted as the maximum transfer rate possible without causing congestion. By deploying the methods studied in this thesis the available bandwidth can be measured without previous knowledge of the network topology. When an estimate of the available bandwidth is obtained, the transfer rate when sending messages between computers can be set to the measured value in order to avoid congestion.

In the thesis an active end-to-end available bandwidth measurement method called "Bandwidth Available in Real Time" (BART for short) is evaluated. BART measures the available bandwidth by injecting probe packets into the network at a given rate and then analysing how this rate has changed on the receiving side. A Kalman filter is used to update the current estimate of the available bandwidth using the new measurement sample.

The focus of the thesis is on how methods, such as BART, function in wireless 802.11 networks, which are very popular in work as well as in home environments. Wireless networks have a different construction compared to many other types of networks and this can affect the accuracy of the measurement methods discussed in this thesis. The effects must be analyzed and understood in order to obtain accurate available bandwidth estimates. Since wireless links are often parts of the network path between a sender and a receiver on the Internet, it is important to study how these links affect the estimates of the available bandwidth.

ISSN 1651-4238

ISBN 978-91-85485-41-3

To my family

Acknowledgements

I would like to thank my supervisor Mats Björkman for seeing potential in me and for hiring me as his PhD student. Mats, you have been my great guide and companion during both private and work related ups and downs. I respect and admire you very much for that!

I would like to thank my supervisor Bob Melander. Without your guidance and support I would never have come this far. You have taught me that everything is possible if you just try hard enough!

I owe my deepest gratitude to Svante Ekelin, Christofer Flinta, Erik Hartikainen, Jan-Erik Mångs, Martin Nilsson and my two supervisors, all involved in the EvaluNet project group. As you know, our team work has been very rewarding for me, and I must admit that it is difficult to find occasions I have been happier at work than after interesting and challenging discussions with you!

I would also like to thank Jonas Neander and Ewa Hansen, both part of the data communication group at Mälardalen University, for their support. You have really enriched my time as a PhD student! I also want to send my great thanks to the rest of the personal at my department for interesting discussions, coffee break laughs and in general for making my department a creative and fun place to work at.

I wish to thank the CoRe-group at Uppsala University for giving me the opportunity of using their facilities when working in Uppsala. I particularly want to thank Christian Rohner for giving me feedback on my PhD thesis and on papers I have written earlier. Also, input and comments on papers from Henrik Lundgren has been rewarding!

I would also like to thank Annikki Welin and Rikard Holm for creative discussions and for helping me with many of the measurements performed during the first year of my PhD studies.

To my mom Elly and dad Lennart, thank you for providing me with chemistry boxes, astronomical telescopes, computers and other fun stuff, it really pointed out the direction of my future! My parents as well as my brother Mattias and sister Jenny have always supported me and encouraged me during my PhD studies. Thank you!

Finally, my deepest thanks goes to Maaret for always motivating and believing in me and making my life outside work enjoyable.

I also wish to thank my sponsors. The Swedish Research Council gave me economical support during my first three years as a PhD student. Then Vinnova and the KK Foundation supported my work through the EvaluNet project. Financial support for travels and alike was granted from ARTES++.

Andreas Johnsson
Västerås, April, 2007

Contents

II Included Papers 35

List of Publications

The following articles are included in this thesis:

A. *On the Analysis of Packet-Train Probing Schemes*, Andreas Johnsson, Bob Melander and Mats Björkman, In proceedings of the International Conference on Communication in Computing, Special Session on Network Simulation and Performance Analysis, Las Vegas, June 2004.

B. *An Analysis of Active End-to-end Bandwidth Measurement Methods in Wireless Networks*, Andreas Johnsson, Mats Björkman and Bob Melander, In proceedings of the 4th IEEE/IFIP End-to-end Monitoring Techniques and Services workshop, Vancouver, Canada, April 2006.

C. *Real-time Measurement of End-to-End Available Bandwidth Using Kalman Filtering*, Svante Ekelin, Martin Nilsson, Erik Hartikainen, Andreas Johnsson, Jan-Erik Mångs, Bob Melander and Mats Björkman, In proceedings of the 10th IEEE/IFIP Network Operations and Management Symposium, Vancouver, Canada, April 2006.

D. *Measuring the Impact of Active Probing on TCP*, Andreas Johnsson and Mats Björkman, In proceedings of the SCS International Symposium on Performance Evaluation of Computer and Telecommunication Systems, Calgary, Canada, July 2006.

E. *On measuring the available bandwidth in wireless 802.11b networks*, Andreas Johnsson and Mats Björkman, Not yet published.

Besides the above articles, I have (co-)authored the following peer-reviewed scientific papers:

I. *On the Comparison of Packet-Pair and Packet-Train Measurements*, Andreas Johnsson, In Proceedings of the First Swedish National Computer Networking Workshop, Arlandastad, September 2003.

II. *Bandwidth Measurements from a Consumer Perspective - A Measurement Infrastructure in Sweden*, Mats Björkman, Andreas Johnsson and Bob Melander, Presented at the Bandwidth Estimation (BEst) Workshop, San Diego, December 2003.

III. *A Study of Dispersion-based Measurement Methods in IEEE 802.11 Adhoc Networks*, Andreas Johnsson, Mats Björkman and Bob Melander, In proceedings of the International Conference on Communication in Computing, Special Session on Network Simulation and Performance Analysis, Las Vegas, June 2004.

IV. *DietTopp: A First Implementation and Evaluation of a Simplified Bandwidth Measurement Method*, Andreas Johnsson, Bob Melander and Mats Björkman, In Proceedings of the Second Swedish National Computer Networking Workshop, Karlstad, November 2004.

V. *Bandwidth Measurement in Wired and Wireless Networks*, Andreas Johnsson, Bob Melander and Mats Björkman, In Proceedings of the Fourth Annual Mediterranean Ad Hoc Networking Workshop, Porquerolles, France, June, 2005.

VI. *How does Available Bandwidth Measurement Methods Affect TCP?*, Andreas Johnsson, In Proceedings of the third Swedish National Computer Networking Workshop, Halmstad, November 2005.

I

Thesis

Chapter 1

Introduction

1.1 Background and motivation

Internet has gained much popularity among the public since the mid 1990's and is now an integrated part of our society. A large range of high-speed broadband providers and the development of new and more efficient Internet applications increase the possibilities to watch movies and live TV, use IP-telephony and share files over the Internet. Such applications demand high data transmission rates, which in turn consume network bandwidth. Since several users must share the common bandwidth capacity on the Internet, there will be locations in the network where the demand is higher than the capacity. This causes network congestion, which has negative impact on both the data transmission rate and transmission quality.

If the status of the network path can be predicted, it is possible to use that information in network applications. For example, the network status can be used to adapt the traffic load in order to avoid congestion on a network path. The network status information can also be used when making a decision on whether to setup a new connection or not. For example, an IP-telephony session requires that the latency and jitter between the two communicating nodes are low. The IP-telephony application also requires that there is enough available bandwidth to transfer the encoded voice data over the network. If these requirements are fulfilled, the connection can be established and the involved people can start talking over the Internet.

The status of a network such as the Internet must be measured. This is because the Internet is a best-effort network that in the general case does not give users opportunities to make bandwidth reservations for their traffic.

This thesis concerns methods for actively measuring the status of the network path between two nodes. One important metric to measure is the end-to-end available bandwidth of a network path. The available bandwidth is in this thesis interpreted as the maximum transfer rate that can be used without causing congestion on any of the links between two nodes. The available bandwidth measurement methods studied in this thesis inject so called probe packets into the network path between the sender and the receiver. By conducting analysis on the separation of the probe packets on the sender and on the receiver side, an estimate of the available bandwidth is obtained.

The main contributions of this thesis are: a model for measuring the available bandwidth in networks where the bottleneck is a wireless 802.11 link. Further, a measurement method called BART is evaluated in both wired and wireless networks. It is also studied how the injected probe packets affect other network traffic.

An alternative approach to measure the available bandwidth is to deploy passive measurement methods. As opposed to active measurement methods, passive measurement methods require access to routers in the network and can only be used by network administrators and other persons with authorized access to the network. Passive measurement methods listen to the network traffic flowing through the network, rather than injecting probe packets. Passive measurement methods are not within the scope of this thesis.

There exist many examples of applications that would benefit from having an updated view of the available bandwidth of the network path. For example, the available bandwidth estimate can be used as input to the adaptive machinery of streaming audio and video. If the available bandwidth is low, the streaming application can reduce its send rate by decreasing the quality of the audio or video, then less data needs to be sent per time unit. On the other hand, if the available bandwidth is high, more data can be transmitted without causing congestion and hence the quality of the audio and video can be increased.

Another example is the following; suppose a user wants to download a large file from a website on the Internet. By deploying active end-to-end measurements he or she (or the web browser itself) can decide which mirror site to use by estimating the available bandwidth to each site and then calculate the minimum download time. Other examples of applications include network diagnosis, load balancing in routers and verification of service level agreements in user-corporation or corporation-corporation relations.

1.2 Outline of the thesis

This thesis is divided into two parts, I and II. Part I provides an overview of the research area, the contributions of this thesis, related work, conclusions and future work. The second part of the thesis consists of five scientific papers labeled A, B, C, D and E, respectively. Papers A - D are peer-reviewed while paper E is not yet published.

1.3 Research area

This thesis focuses on how to actively measure end-to-end performance in IP networks. Most methods discussed in this thesis require a probe-packet sender and a receiver. The sender injects so called probe packets into the network path with a pre-defined packet separation, corresponding to a probe rate. The probe packets are affected by the current status of the network, and this may change the probe-packet separation. This separation is measured at the receiver by time stamping each probe-packet arrival. Depending on the analysis method, different performance metrics can be estimated, such as the available bandwidth.

In this section the research area of active end-to-end measurement methods is introduced. First the network model assumed is briefly described. Then a definition of link capacity, available bandwidth and cross traffic is given. A high level description on how to measure the available bandwidth is presented. More detailed descriptions of methods and tools for measuring available bandwidth are given in Chapter 3. This section ends with a review of research challenges addressed in this thesis.

1.3.1 Network model

In this section the network model is presented. Even though it is simplistic, encompassing few parameters, it is shown in Part II of this thesis that the available bandwidth measurement methods studied produce accurate estimates in both testbed and Internet scenarios.

The networks that the methods studied in this thesis intend to measure are best-effort packet-switched network. This means for example that no bandwidth reservations can be made and that all network traffic is transmitted as discrete packets.

A network path consists of several consecutive links with different capacities C. The links in the network are inter-connected by routers and each router has several incoming and outgoing links. The task of the router is to forward incoming traffic to an outgoing link that leads towards the destination of the data packets.

A router is modeled as a set of queues, one for every outgoing link. If the amount of network traffic that is forwarded to a specific outgoing link is larger than the capacity of that link, the data packets are typically stored in a first-in-first-out fashion. This leads to a delay as a result of the congestion. In case the buffered traffic exceeds the size of the queue data packets are dropped. This situation is called persistent congestion.

1.3.2 Link capacity, available bandwidth and cross traffic

The term *available bandwidth* has had varying definitions in the literature. In this section, definitions of both *link capacity* and *available bandwidth* as well as a short description of *cross traffic* are given with the thesis topic in mind, which is the available bandwidth and link capacity at the IP layer in IP networks. The definitions of link capacity and available bandwidth are taken from [1].

Each link j in an end-to-end network path has a certain link capacity C_j, which is defined as the maximum transmission rate over that link (bit/s). The maximum rate is determined by the link interfaces used by the nodes in the network path. The link capacity is typically constant, but can shift in e.g. wireless 802.11 networks when the radio quality reaches some upper or lower threshold. The link utilization for each link j varies over time due to cross traffic[1], as illustrated in Figure 1.1. A variety of parameters can be used to describe the cross traffic, such as the intensity distribution, the packet size distribution and the aggregation of cross-traffic flows. Measurement of these parameters in different scenarios and networks is a research area in itself (e.g. in [2] the packet-size distribution was studied).

The intuitive definition of available bandwidth is the unused portion of the link capacity during some time interval. The portion of the link capacity that is in use during a measurement corresponds to the cross traffic $X_j = X_j(t, \Delta)$, which is defined by

$$X_j = \frac{1}{\Delta} A_j[t - \Delta, t] \tag{1.1}$$

---

[1] The cross traffic is packets being transmitted over a network link.

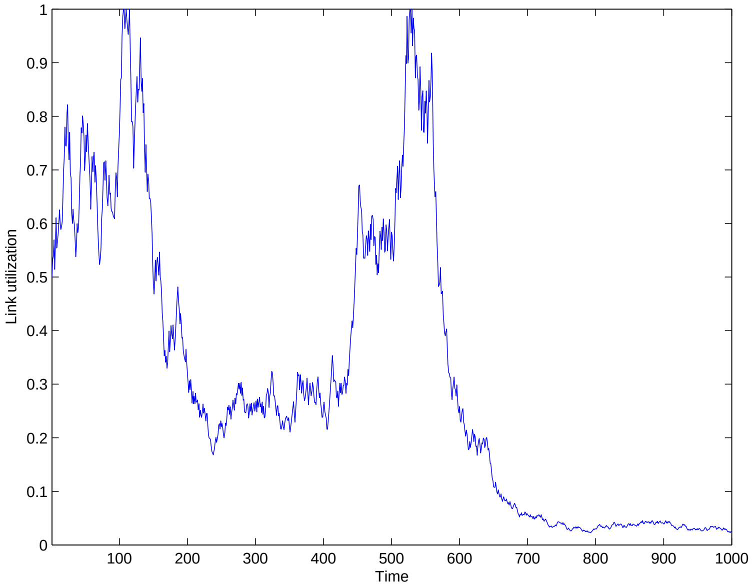

Figure 1.1: Link utilization as a function of time.

where $A_j[t-\Delta, t]$ is the number of cross traffic bits over link j during a time interval Δ. The available bandwidth $B_j = B_j(t, \Delta)$ is defined as

$$B_j = C_j - X_j. \tag{1.2}$$

The available bandwidth of an end-to-end path is then defined as the minimum available bandwidth for each link j. That is,

$$B = \min_j B_j. \tag{1.3}$$

The link corresponding to B is often referred to as the bottleneck link or the tight link of the network path.

This definition of available bandwidth can be interpreted as follows; if sending packets at a rate below B, the packets will not cause congestion on the network path. On the other hand, if sending packets at a higher rate than B, the packets will be queued at the bottleneck link, and thus induce congestion.

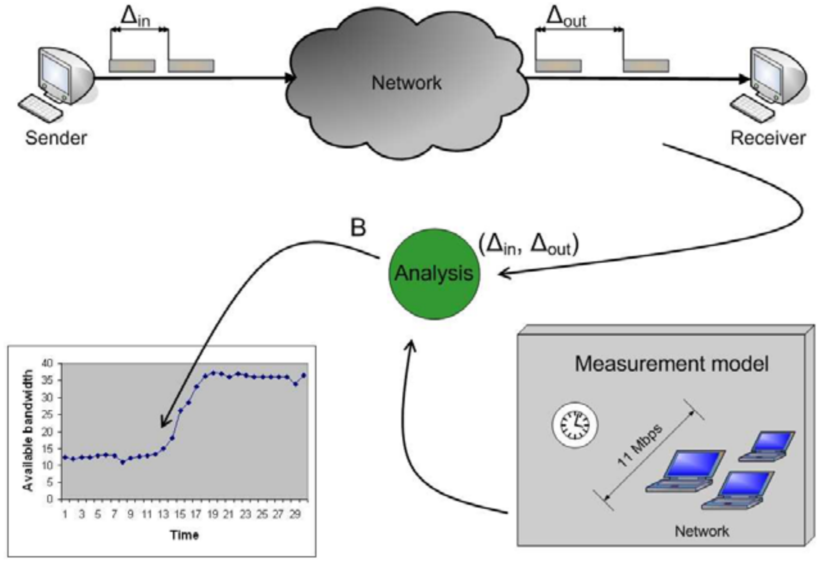

Figure 1.2: High-level description on how to actively measure the end-to-end available bandwidth.

1.3.3 Measuring the available bandwidth

Assuming that two nodes want to measure the available bandwidth on the end-to-end path linking them together, then so called *active probing* is a feasible approach. This section describes, at an abstract level, how the end-to-end available bandwidth can be measured using methods that rely on active probing.

The principle of active probing is to inject so called *probe packets* into the path to be measured. Assume that the sender in Figure 1.2 wants to measure the available bandwidth of the path to the receiver. Then the sender injects probe packets with initial mean packet separation Δ_{in}. The probe packets traverse the network path, and along the way the separation may change due to cross traffic or link properties, such as the capacity of the links. At the receiver side each probe packet is time stamped when it arrives. Then Δ_{out} can be calculated, which is the separation between probe packets at the receiver side. If the mean input separation of several successive measurements is equal to the output separation, it can be assumed that no congestion has occurred on the path. That is, the probe rate (proportional to the inverse of the probe-packet separation) is less than the available bandwidth B. On the other hand, if the mean output separation is greater than the input separation, the probe packets have experienced congestion and thus the separation has increased in proportion to the amount of cross traffic on the path. The probe rate is in this case above the available bandwidth. Most available bandwidth measurement methods rely on this idea, which is called *self-induced congestion*. It originates from the fact that the probe packets themselves induce congestion on the path for a very short time interval in order to be able to measure how much cross traffic there actually is, and then calculate the available bandwidth.

Both Δ_{out} and Δ_{in} are fed into an analysis algorithm. The analysis is based on a measurement model that is used to produce estimates of the available bandwidth using samples of Δ_{out} and Δ_{in} as input. (Different measurement models are described in Chapter 3.) The resulting available bandwidth estimate can then be used in an application such as a network monitoring tool that monitors an end-to-end path and plots the estimates in a graph, as illustrated in Figure 1.2.

The approach of self-induced congestion requires no previous knowledge of the network topology. This is desirable since the link property information is hard or impossible to retrieve for the average user on the Internet.

In Chapter 3 a more comprehensive description of different methods and tools in this area of research is presented. The scheme of how to inject the probe packets and the analysis of the probe-packet time stamps vary depending on the method.

1.3.4 Research challenges

The research challenges addressed in this thesis are summarized below.

- Find and evaluate new and lightweight active end-to-end measurement methods that are able to obtain bandwidth estimates in reasonable time without consuming too much resources of the network or on the probe-packet sender and receiver nodes. It is also important that new methods minimize the time between the actual measurement and presentation of the bandwidth estimate to the application. Old estimates are outdated since the network conditions may have changed.

- It is important to understand whether bandwidth estimation methods report correct values in wireless networks, such as 802.11 or 3G. For many of the bandwidth estimation methods it is assumed that the forwarding and queuing mechanism in the routers is first-in-first-out based. What happens in the case where packets are assigned priority based on source or destination address or application type? Do bandwidth estimation methods still report a valid estimate?

- It is important to study how the probe packets injected by the available bandwidth measurement methods that exist today affect other traffic flows on the network. Is there a difference between how the probe packets affect UDP traffic flows compared to flows mainly consisting of TCP? Is it possible to inject the probe packets in a way that minimizes the impact on other traffic flows while at the same time keeps the accuracy of the measurement samples?

Even though the active end-to-end bandwidth measurement research area has matured over recent years there still exist many important research challenges to be addressed by researchers in the future. These are presented in Chapter 4.

1.4 Research methodology

To produce the research results presented in this thesis two main research methodologies have been used. Analytical methodology has been the primary tool when modeling dependencies while experimental methodology has been applied when analytical verification is infeasible due to the complexity of the problem.

The primary task of experimental research methodology is to study the relation between cause and effect. A set of hypotheses regarding a system or a phenomenon are tested by conducting experiments. Thereafter conclusions are drawn which in turn will lead to new hypotheses or input to the analytical methodology. In experimental research the effect of single variables are isolated and thus the impact of that variable can be understood and perhaps modeled. In the available bandwidth measurement research area it is for example interesting to study the impact of the cross-traffic distribution on the available bandwidth estimate produced by different methods.

The experiments presented and analyzed in this thesis have been performed over real Internet paths, in testbed laboratories or in network simulation tools. Experimental results obtained from laboratories and simulation tools are possible to verify since all variables are known. However, experiments performed over real Internet paths are harder to verify due to the fact that it is a non-trivial task for researchers to gain access to the core Internet.

Another aspect of experimental research methodology is the reproducibility of the study. By conducting experiments in network simulations and in testbeds, and then documenting all parameters, the experiments can be repeated with similar results. Experiments in real networks, such as the Internet, are harder to reproduce since it is for example impossible to control and re-produce the cross-traffic conditions at a given time.

1.5 My contributions

The main scientific contributions of the work presented in this thesis, conducted by myself and in collaboration with others, are the following:

- A model describing how bandwidth measurement methods are affected by the properties of wireless 802.11 networks. The model is an extension of the TOPP model [3] (an available bandwidth measurement model). The extended model describes the dependencies between available bandwidth, link capacity estimates, probe-packet size and cross-traffic rate.
- The parameters in the Kalman filter used by BART [1] (a bandwidth measurement method) have been tuned to increase the accuracy of the estimates, when the bottleneck is a wireless 802.11 link. The parameter configuration is found to be different when measuring in wireless networks compared to wired networks. Further, the BART method has also been evaluated in both wired testbed and Internet scenarios. The results show that BART reports accurate bandwidth estimates.
- A study on how different active end-to-end available bandwidth measurement methods affect other traffic flows. In a simulation based study it is shown that the impact of probe packets on TCP flows is limited even though the probe packets may increase the round-trip time and the packet loss experienced by the TCP flows.

Other results include:

- ○ A framework describing the probe-packet train and cross-traffic packet interactions at the discrete packet level has been developed.
- ○ Identification of problems concerning the send mechanism in Linux of probe packets when injected into a wireless 802.11b network.

- ○ Implementation of two bandwidth measurement methods, DietTOPP and BART, in C/C++ for UNIX based systems.
- ○ Construction of a testbed network consisting of both wired single-access and wireless multi-access links.

Chapter 2

Summary of the papers

The second part of this thesis is a collection of five papers; paper A, B, C, D and E. Papers A - D have been published in peer-reviewed conference or workshop proceedings while paper E has not yet been published. This chapter gives a short summary of each paper along with editorial comments and comments on my contribution. The editorial comments are intended as correction of errors and to simplify for the reader of the papers.

Papers A and D focus on the interaction between probe packets and cross-traffic packets. Paper B reports on a study on how wireless 802.11 bottleneck link properties affect available bandwidth estimates. Paper C presents a bandwidth measurement method named Bandwidth Available in Real Time (BART), along with an evaluation of the method in various settings. Paper E is a continuation of papers B and C. The foci of paper E are on general problems concerning measurements of available bandwidth in wireless networks and on tuning of parameters in BART in order to increase accuracy of the estimates, when the bottleneck is a wireless 802.11 link.

2.1 Paper A: On the Analysis of Packet-Train Probing Schemes

Andreas Johnsson, Bob Melander and Mats Björkman, ”On the Analysis of Packet-Train Probing Schemes”, In proceedings to the International Conference on Communication in Computing, Special Session on Network Simulation and Performance Analysis, Las Vegas, USA, June 2004.

Summary: This paper describes probe packet and cross-traffic packet interactions at the discrete packet level. Three main interaction-types are identified; mirror patterns, chain patterns and quantification patterns, respectively. Using the pattern terminology, the difference between using the statistical mean and median operators to reduce noise in the measurements are explored.

My contribution: I have written most of the paper, co-developed the measurement tool and performed all experiments.

2.2 Paper B: An Analysis of Active End-to-end Bandwidth Measurements in Wireless Networks

Andreas Johnsson, Mats Björkman and Bob Melander, ”An Analysis of Active End-to-end Bandwidth Measurements in Wireless Networks”, In proceedings to the 4th IEEE/IFIP End-to-end Monitoring Techniques and Services workshop, Vancouver, Canada, April 2006.

Summary: In this paper a study on how the properties of wireless 802.11 bottleneck links affect bandwidth estimates (i.e. both available bandwidth and link capacity) produced by methods related to the TOPP model is presented. The results from experiments performed in a testbed environment illustrate that the probe-packet size is crucial to the estimate of the available bandwidth. The estimates decrease with decreasing probe-packet size. From the experimental results it is also evident that the estimated link capacity, using the TOPP model, is dependent on the probe-packet size and the cross-traffic rate. The reason for the dependency between bandwidth estimates, probe-packet size and cross-traffic rate is due to the link-layer mechanisms in 802.11 networks. This observation was developed into an extension of the TOPP model that describes the dependencies in detail.

Editorial comments: In this paper there is no differentiation made between cross traffic sent by other nodes in the wireless network and cross traffic injected by the probe-packet sender (which usually is zero during a measurement session). This should however not affect the overall understanding of the extended TOPP model. The cross traffic that affect the function $T_k(x)$, in Equation 6.8 in the paper, originates from other nodes using the same wireless link while the ”single” x corresponds to the cross traffic generated at the node injecting the probe packets. In paper E this is noted.

My contribution: I have written most of the paper, co-analyzed the experimental results, conducted all measurements and implemented the measurement tool DietTopp.

2.3 Paper C: Real-time Measurement of End-to-End Available Bandwidth Using Kalman Filtering

Svante Ekelin, Martin Nilsson, Erik Hartikainen, Andreas Johnsson, Jan-Erik Mångs, Bob Melander and Mats Björkman, ”Real-time Measurement of End-to-End Available Bandwidth Using Kalman Filtering”, In proceedings to the 10th IEEE/IFIP Network Operations and Management Symposium, Vancouver, Canada, April 2006.

Summary: This paper presents the BART method for measuring and estimating the system state of an end-to-end path. From the system state both the available bandwidth and the link capacity of the bottleneck link can be derived. BART, which is an offspring of the TOPP method, makes use of a Kalman filter in order to track system state changes in real time. The measurement model BART makes use of is piecewise linear. Since the Kalman filters require a linear model, it is also discussed in the paper how to overcome this limitation.

In the paper an evaluation of the BART method in testbed networks as well as over real Internet paths is presented. The evaluation shows that BART is very accurate in most cases. Since BART estimates the available bandwidth in real time, results from experiments where the available bandwidth varies are also discussed. BART reacts quickly to the changes in available bandwidth. The accuracy of BART is also compared to results obtained by pathChirp. The impact of the probe-packet size and the probe-packet train length is briefly studied. Large probe packets in combination with long probe-packet trains is the best choice in order to get accurate results when performing the experiments discussed in this paper.

Editorial comments: BART make use of the TOPP model, which describes the relation between the offered probe rate o (probe rate at the sender side) and the measured probe rate m (probe rate measured at the receiver). In this paper the variable names has changed to u for the offered rate and r for the measured rate at the receiver.

My contribution: I have written parts of the evaluation section, participated in discussions on how to evaluate BART, co-designed evaluation scenarios and participated in conducting measurements. I have also implemented the measurement method into a tool used in the evaluation process.

2.4 Paper D: Measuring the Impact of Active Probing on TCP

Andreas Johnsson and Mats Björkman, ”Measuring the Impact of Active Probing on TCP”, In proceedings to the SCS International Symposium on Performance Evaluation of Computer and Telecommunication Systems (SPECTS’06), Calgary, Canada, July 2006.

Summary: This paper discusses the impact of active probing, such as the impact of probe packets injected by available bandwidth measurement methods, on TCP. In the paper, simulation results show that the available bandwidth on a link in an end-to-end path actually can increase when injecting probe packets into a path even though the probe packets add traffic to the network, if the cross traffic consists of TCP flows. This is because the probe-packet trains may cause packet losses that in turn trigger the TCP backoff mechanism. Thus, the overall TCP throughput is reduced. To investigate whether it is possible to lower the impact of the probe packets on the TCP flows, different probe-packet train lengths were used when measuring the available bandwidth in order to study the impact on the TCP flows. The number of samples, that is the number of probe-packet separations per time unit, was however constant. The simulation results show that probe-packet trains of ”medium length” is preferable. Short probe-packet trains increase the TCP round-trip time estimate while long probe-packet trains increase TCP loss. Both have negative effect on TCP throughput.

In the paper it is noted that the overall impact of the probe packets on TCP throughput is rather low.

My contribution: I have written most of the paper, co-analyzed the simulation results, constructed the simulation setup and performed all simulations.

2.5 Paper E: On measuring the available bandwidth in wireless 802.11b networks

Andreas Johnsson and Mats Björkman, ”On measuring the available bandwidth in wireless 802.11b networks”, Not yet published.

Summary: In this paper it is further studied how the properties of wireless 802.11 networks affect bandwidth measurement methods. Two problems related to the operating system drivers are identified. These problems can bias the measurement accuracy of available bandwidth and link capacity. In the paper the accuracy of the BART method is also investigated, when measuring in wireless 802.11 networks.

It is possible to tune the Kalman filter used by BART and thereby to give different weight to the new measurement sample on the current estimate of the available bandwidth. In this paper the BART method is tuned in order to increase the accuracy when measuring in networks where the bottleneck is a wireless 802.11 network. It is shown that the configuration of the tunable parameters is different from the configuration used in wired single-access networks.

My contribution: I have written most of the paper, co-analyzed the experimental results and performed all measurements.

Chapter 3

Related work

Much work has been published in the active end-to-end bandwidth estimation research area in recent years, and the area has started to mature. Many different available bandwidth estimation methods and tools have been developed and each new method tries to address new aspects of measuring available bandwidth. Hence, each new method contributes to the common understanding on how to perform active end-to-end available bandwidth measurements.

The problem of measuring the end-to-end available bandwidth is complex because there are many theoretical issues as well as real-world problems that must be considered, such as the impact of the cross-traffic distribution on the probe packets, or clock tick resolution issues on the sender and receiver side. The methods differ in the design, how the probe packets are injected and how the analysis is performed, and thus makes different assumptions and address different challenges regarding active end-to-end measurements. This leads to the conclusion that there may not be one available bandwidth measurement method that is optimal in all scenarios. Rather, the method to be applied depends on the demands of the application using the measurement method or on what is known about the characteristics of the end-to-end network path.

This chapter reviews important literature related to this thesis. Several methods have been developed for actively measuring the end-to-end available bandwidth. These are mentioned in Section 3.1. In Section 3.2 work of more theoretical nature is presented.

3.1 Methods and tools

Well known active end-to-end bandwidth measurement methods are for example Pathchirp [4], Pathload [5], Spruce [6] and TOPP [3]. Among newly developed methods, ABget [7] and BART [1] can be mentioned. The above methods can be classified in two categories: *direct probing* and *iterative probing* [8].

In direct probing, each injection of probe packets (as pairs or in longer sequences as trains) gives a sample of the available bandwidth. *Spruce* is one such method. As other methods that rely on direct probing, Spruce needs knowledge of the bottleneck link capacity to obtain estimates of the available bandwidth. By comparing the input rate and the output rate, the cross-traffic rate can be measured. Then, if the link capacity is known the available bandwidth is calculated as the link capacity minus the cross-traffic rate. The advantage of direct probing is that one probe-packet train can be enough to determine the available bandwidth.

In iterative probing no knowledge of the end-to-end path is needed. Instead, these methods only rely on self-induced congestion. The idea is the following: inject a set of probe packets (e.g. in pairs or in trains) with a predefined initial probe rate into a network. If the initial probe rate is higher than the available bandwidth, the probe packets will be queued after each other at the bottleneck and cause congestion for a short time interval. In such a case, the mean separation between the probe packets will increase, which is equivalent to a decrease in the received probe rate. The increase in the probe-packet separation is proportional to the amount of cross traffic. However, if the initial probe rate is equal to the received probe rate it is assumed that the packets did not have to queue and thus the end-to-end path is not congested. That is, the initial probe rate is less than the available bandwidth. By sending several probe packets with different input rates and then applying an analysis method, the available bandwidth is located. There exist several methods that exploit this feature; such as ABget, BART, Pathchirp, Pathload and TOPP.

In the following paragraphs, descriptions of a selection of active end-to-end bandwidth estimation methods are given. A more in-depth review of available bandwidth measurement theory, methods and tools can be found in [9]. In [10] the problems of comparing different available bandwidth measurement methods are discussed. Further, in [11] a comparison of public end-to-end available bandwidth estimation methods on high-speed links is presented.

Spruce Spruce [6] is an example of direct probing, which assumes that the capacity of the bottleneck is known. Spruce injects probe packets with sepa-

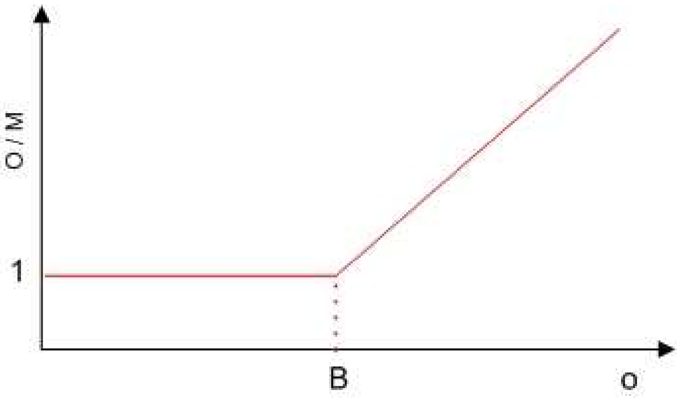

Figure 3.1: The rate-response curve obtained by TOPP.

ration Δ_{in} and measures the separation Δ_{out} on the receiver side. Then the rate of the cross traffic during the network path traversal of the probe packets is $\frac{\Delta_{out}-\Delta_{in}}{\Delta_{in}}C$ where C is the capacity of the bottleneck link. Then the available bandwidth B is

$$B = \Big(1 - \frac{\Delta_{out} - \Delta_{in}}{\Delta_{in}}\Big)C. \tag{3.1}$$

In [12] it was shown that Spruce and other methods that rely on direct probing tend to underestimate the available bandwidth in the general case.

TOPP TOPP [3] injects probe packets at rate o in an interval o_{min} to o_{max} and measures the received rate m at the receiver node. If plotting o/m versus o the rate-response curve is obtained, see Figure 3.1. If $o/m = 1$ then the probe packets did not cause congestion on the network path. Otherwise, if $o/m > 1$ congestion has been induced into the network path and the rate-response curve deviates from $o/m = 1$ in a linear fashion.

After the measurement session TOPP deploys linear regression in order to find the equation of the sloping segment. The available bandwidth is defined as the intersection of the line $o/m = 1$ and the sloping segment, seen in Figure 3.1. That is, the point corresponding to an offered rate that just not congest the network path. Having the equation of the sloping segment, the available bandwidth is easily found. TOPP is explained in more detail in Part II of this thesis.

Pathload Pathload [5] is an implementation of the Self-Loading Periodic Streams (SLoPS) methodology [5][13]. Using this methodology the end-to-end available bandwidth is measured. The basic idea of SLoPS is explained below.

A sender transmits a packet train with a pre-defined probe rate to a receiver. The sender time stamps the send time and the receiver time stamps each probe packet when it arrives. The time stamp difference, defined as the *one way delay* (OWD) is calculated for each packet ($D_1, D_2, ..., D_K$).

If the initial rate is higher than the available bandwidth, the length of the router queue grows when additional probe packets in the train are received to the router. Due to this fact, the OWD values will have an increasing trend (i.e. $D_K > D_{K-1} > ... > D_1$). When the initial probe rate is less than the available bandwidth, the router queues will not grow as a result of the injection of the probe train. Hence, the OWD will be more or less stable. Statistical tests are used to determine whether the OWD values are stable or increasing.

By sending probe trains at different pre-defined probe rates, the available bandwidth is found by binary search. If the initial probe rate causes an increasing OWD trend, the probe rate is above the available bandwidth. The probe-train rate for the next train is then decreased. On the other hand, if the probe train does not cause an increasing OWD trend, the initial probe rate is less than the available bandwidth. In this case the probe rate is increased. This process is iterated until a satisfactory accuracy of the available bandwidth is obtained.

Pathchirp Pathchirp [4] is another tool that measures the end-to-end available bandwidth and exploits the iterative method. Instead of sending probe-packet trains it sends *chirps* of probe packets. A chirp is essentially a packet train, but the separation between the probe packets are exponentially decreasing, $d_n = T\gamma^n, d_{(n-1)} = T\gamma^{(n-1)}, ..., d_{(2)} = T\gamma^2, d_1 = T\gamma^1$, where d_i is the separation and T and γ are constants. By using chirps the end-to-end path is probed at different initial probe rates using just one chirp.

In the analysis the probe packet send time is compared to the queuing delay that arises when the initial probe rate is above the available bandwidth. The queuing delays are derived from comparing probe packet send and receive times. In Figure 3.2 an example is shown. It illustrates the measurement results obtained from one chirp. When the queuing delay is zero, the probe packets were transmitted without queuing. During the excursions, shown in the figure, the cross traffic is more intensive and hence causes queuing. The first few excursions go back to zero because the cross traffic is bursty (i.e. sometimes the

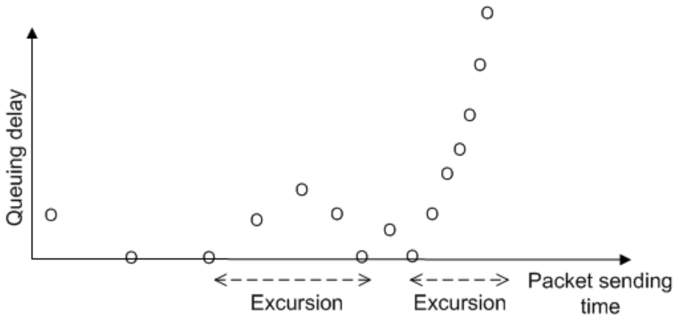

Figure 3.2: Packet sending time versus queuing delay by Pathchirp.

cross traffic is absent during parts of a chirp). The last excursion in the figure does not return to zero. This is when the send probe rate has exceeded the bottleneck link capacity.

By analyzing the excursions from one chirp it is possible to estimate the end-to-end available bandwidth. By injecting several chirps over time, changes in the available bandwidth can be tracked.

BART BART [1] injects probe packets at random rates in the interval u_{min} to u_{max}. A system state vector describes the sloping segment of the rate-response curve, similar to the TOPP model, see Figure 3.3. In BART the system state is updated for each measurement sample using a Kalman filter. This way an estimate of the available bandwidth is obtained in real time. The Kalman filter tells how much to trust the new sample compared to the current estimate of the available bandwidth.

The update of the system state is illustrated in Figure 3.3. The system state describes the solid line which results in an available bandwidth estimate B (compare to the TOPP method). A new sample is obtained by injecting probe packets into the network path. The system state changes due to the new sample, and a new available bandwidth estimate B_x is produced. BART is described in more detail in the second part of this thesis.

ABget ABget [7] relies on the same type of analysis as Pathload. The novel thing about ABget is that it does not require a probe-packet sender and receiver. Instead it makes use of a fake TCP client which instructs an ordinary TCP server to send packets according to a scheme similar to sending probe-packet

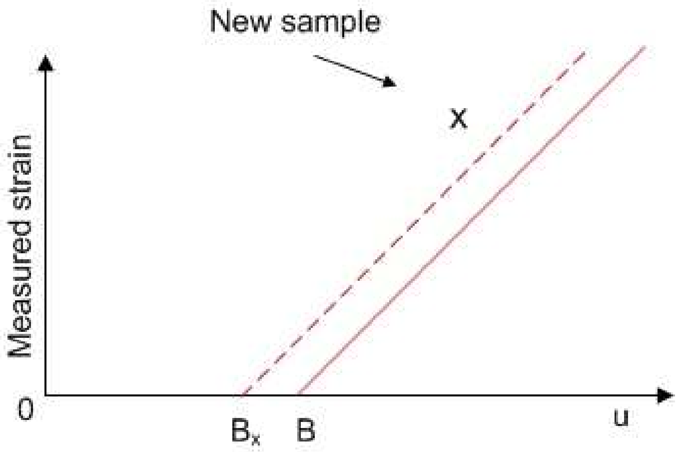

Figure 3.3: BART.

trains. The actual analysis of the received probe packets on the client side is done using the ideas of increasing one-way delay during congestion, taken from Pathload.

3.1.1 Link capacity estimation

There also exist methods for estimating the link capacity (not just the available bandwidth) of the bottleneck link. Pathrate [14], TOPP and BART are three examples. Pathrate looks for a capacity mode obtained from packet-pair probing while TOPP and BART produce the link capacity as a side effect when estimating the available bandwidth.

Pathneck [15] is a tool that uses ICMP packets and ordinary probe packets to find the actual location of a bottleneck link in an end-to-end path. However, it cannot estimate the link capacity or the available bandwidth.

There exist a variety of older tools that for example uses ICMP and variable packet sizes in order to estimate the link capacity, but these methods fail in many cases [16].

3.2 Theoretical work

Much theoretical research has been produced throughout the years. In this subsection a short review of the literature is given.

In [14] it was studied how probe-packet pairs and trains are affected by cross traffic. Depending on the scheme used, different modes are visible in histograms of probe-packet time separations. In this study the authors defined one of the modes as the asymptotic dispersion rate, which was called the proportional share in [17]. This value corresponds to the proportion of the link capacity in use by the probe packets. In [14] a study of the impact of probe packet size was also investigated. The findings lead forward to the link capacity estimation method Pathrate.

In [17][3] discussion about the proportional share, among other things, lead forward to the TOPP model to measure the end-to-end available bandwidth and link capacity.

In [18] a delay-variation model for packet-pair like methods was developed. Using this model several histogram signatures could be identified. The modes in the histograms corresponded to bottleneck link and secondary bottleneck link capacities. A secondary bottleneck link is a link that is congested when further increasing the probe-packet rate, even though the first bottleneck link is congested. The impact of the probe-packet size to link capacity estimation methods was also discussed.

The effect of layer-2 store-and-forward devices on per-hop capacity estimation was studied in [16]. In this paper it was shown why many older methods for estimating bottleneck link capacity reported erroneous estimates.

A description on how the one-way delay within a probe-packet train varies depending on different parameters such as cross traffic and train length was presented in [13]. By investigating the one-way delay the authors could conclude that the one-way delay increased within a probe-packet train if the send rate was above the available bandwidth. Otherwise the one-way delay remained constant. A method for detecting increasing trends was developed and resulted in the tool Pathload for estimating available bandwidth.

In [19][20][21] the rate response curve (seen in e.g. Figure 3.1) was investigated in-depth. In this work a more comprehensive mathematical model for describing the cross-traffic effect on the rate response curve was presented. However, it has been shown by experiments that simpler models (used by e.g. TOPP and BART) for describing the rate response curve give good estimates of the available bandwidth as well [1][22].

Sometimes the available bandwidth itself is not sufficient for an application. In [23] a study of how to measure the available bandwidth variation range is presented. The available bandwidth is a mean value over some time and the variation range describes how much the cross traffic fluctuates during this interval. The variation range is defined using second order statistics such as the variance. That method was implemented in a tool called Pathvar.

Initial work on how broadband access links, such as 802.11 and ADSL, affect bandwidth measurements obtained from different measurement methods was presented in [24]. Here they showed how the probe-packet size affects the available bandwidth estimates obtained from a set of tools.

The above literature survey, together with the presentation of a selection of bandwidth measurement methods, covers much of the important research made in this area. The papers included in part two of this thesis further develop the research area as described in Section 1.5.

Chapter 4

Conclusions and future research challenges

In this thesis methods for actively measuring the end-to-end available bandwidth have been studied. An evaluation of BART has been presented along with a detailed study of how 802.11 link properties affect available bandwidth estimates produced by methods such as BART and TOPP. The estimates are dependent on the probe-packet size, contrary to what is seen in wired single-access networks.

Further, the Kalman filter deployed by BART has been tuned to increase the accuracy of link capacity and available bandwidth estimates when measuring in wireless networks. The parameter configuration was found to be different compared to the one used in single-access wired networks.

In the thesis it has also been studied in detail how the probe packets and the cross-traffic packets interact with each other. For example, it has been investigated how the probe packets affect TCP flows that share a network path. Depending on the probe-packet train length, either the TCP round-trip time or the TCP loss rate increase, which both cause decreased TCP performance. By simulations it was shown that the combined effect was minimized when the probe-packet sender injected ”medium length” probe-packet trains.

This thesis has given answers to some important research questions mentioned in Chapter 1, but many challenges remain. Some future research challenges not addressed in the thesis, or where more research is needed, are summarized in the list below.

- Today it is possible to measure the end-to-end available bandwidth. In the future it is likely that a more complete picture of the network status is needed. How are active network tomography measurement methods supposed to function? Can such methods be used for bandwidth prediction for a single path, including not only bandwidth estimation but also bandwidth prediction?

- For many of the bandwidth estimation methods it is assumed that the forwarding mechanism in the routers is first-in-first-out based. What happens in the case where packets are assigned priorities based on source or destination address or application type? Do bandwidth estimation methods still report valid estimates? It is also important to understand whether bandwidth estimation methods report correct values in wireless networks, such as 802.11 or 3G.

- A next step is also to integrate bandwidth estimation methods with suitable applications such as streaming media tools, transport protocols or server selection algorithms.

- It would be interesting to investigate whether it is possible to send packets related to the application as probe packets. In that way the total load on the network is decreased. For example, can the packets from a streaming media application be injected in a measurement-like fashion?

Bibliography

[1] S. Ekelin, M. Nilsson, E. Hartikainen, A. Johnsson, J. Mångs, B. Melander, and M. Björkman. Real-time measurement of end-to-end available bandwidth using kalman filtering. In *Proceedings of the IEEE/IFIP Network Operations and Management Symposium*, Vancouver, Canada, 2006.

[2] Yongmin Choi, Heung-No Lee, and Anurag Garg. Measurement and analysis of wide area network (wan) traffic. In *Proceedings of the SCS Symposium on Performance Evaluation of Computer and Telecommunication Systems*, Vancouver, Canada, July 2000.

[3] Bob Melander, Mats Björkman, and Per Gunningberg. Regression-based available bandwidth measurements. In *Proceedings of the SCS Symposium on Performance Evaluation of Computer and Telecommunications Systems*, San Diego, USA, July 2002.

[4] V. Ribeiro, R. Riedi, R. Baraniuk, J. Navratil, and L. Cottrel. pathChirp: Efficient available bandwidth estimation for network paths. In *Proceedings of the Passive and Active Measurement Workshop*, San Diego, USA, 2003.

[5] Manish Jain and Constantinos Dovrolis. Pathload: a measurement tool for end-to-end available bandwidth. In *Proceedings of the Passive and Active Measurement Workshop*, Ft Collins, USA, March 2002.

[6] Strauss, Katabi, and Kaashoek. A measurement study of available bandwidth estimation tools. In *Proceedings of the ACM SIGCOMM Internet Measurement Workshop*, Miami, USA, 2003.

[7] Demetres Antoniades, Manos Athanatos, Antonis Padadogiannakis, Evangelos P. Markatos, and Constantine Dovrolis. Available bandwidth measurement as simple as running wget. In *Proceedings of the Passive and active measurement workshop*, Adelaide, Australia, 2006.

[8] Manish Jain and Constantinos Dovrolis. Ten fallacies and pitfalls in end-to-end available bandwidth estimation. In *Proceedings of the ACM SIGCOMM Internet Measurement Conference*, Sicily, Italy, October 2004.

[9] R.S. Prasad, M. Murray, C. Dovrolis, and K. Claffy. Bandwidth estimation: metrics, measurement techniques, and tools. IEEE Network Magazine, 2003.

[10] Frederico Montesino-Pouzols. Comparative analysis of active bandwidth estimation tools. In *Proceedings of the Passive and Active Measurement Workshop*, Antibes Juan-les-Pins, France, 2004.

[11] A. Shriram, M. Murray, Y. Hyun, N. Brownlee, A. Broido, M. Fomenkov, and K. Claffy. Comparison of public end-to-end bandwidth estimation tools on high-speed links. In *Proceedings of the Passive and Active Measurement workshop*, Boston, USA, 2005.

[12] Li Lao, Constantinos Dovrolis, and M.Y. Sanadidi. The probe gap model can underestimate the available bandwidth of multihop paths. In *ACM SIGCOMM Computer Communications Review*, October 2006.

[13] Manish Jain and Constantinos Dovrolis. End-to-end available bandwidth: Measurement methodology, dynamics, and relation with TCP throughput. In *Proceedings of the ACM SIGCOMM conference*, Pittsburg, USA, August 2002.

[14] Constantinos Dovrolis, Parameswaran Ramanathan, and David Moore. What do packet dispersion techniques measure? In *Proceedings of the IEEE Infocom conference*, Anchorage, USA, April 2001.

[15] Ningning Hu, Li (Erran) Li, Zhuoqing Morley Mao, Peter Steenkiste, and Jia Wang. Locating internet bottlenecks: algorithms, measurements, and implications. In *ACM SIGCOMM: Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications*, Portland, USA, 2004.

[16] Ravi S. Prasad, Constantinos Dovrolis, and Bruce A. Mah. The effect of layer-2 store-and-forward devices on per-hop capacity estimation. In *Proceedings of the IEEE Infocom conference,* New York, USA, April 2002.

[17] Bob Melander, Mats Björkman, and Per Gunningberg. A new end-to-end probing and analysis method for estimating bandwidth bottlenecks. In *Proceedings of IEEE Globcom,* San Francisco, November 2000.

[18] Attila Pásztor and Darryl Veitch. The packet size dependence of packet pair like methods. In *Tenth International Workshop on Quality of Service*, Miami Beach, USA, May 2002.

[19] Xiliang Liu, Kaliappa Ravindran, and Dmitri Loguinov. Single-hop probing asymptotics in available bandwidth estimation: Sample-path analysis. In *Proceedings of the ACM SIGCOMM Internet Measurement Conference*, Taormina, Italy, 2004.

[20] Xiliang Liu, Kaliappa Ravindran, and Dmitri Loguinov. Multi-hop probing asymptotics in available bandwidth estimation: Stochastic analysis. In *Proceedings of the ACM SIGCOMM Internet Measurement Conference*, Berkeley, USA, 2005.

[21] Xiliang Liu, Kaliappa Ravindran, and Dmitri Loguinov. Measuring probing response curves over the ron testbed. In *Proceedings of the Passive and Active Measurement workshop*, Rio de Janeiro, Brazil, 2006.

[22] Andreas Johnsson, Bob Melander, and Mats Björkman. Diettopp: A first implementation and evaluation of a new bandwidth measurement tool. In *Proceedings of the Swedish National Computer Networking Workshop*, Karlstad, Sweden, 2004.

[23] Manish Jain and Constantinos Dovrolis. End-to-end estimation of the available bandwidth variation range. In *Proceedings of the ACM SIGMETRICS conference*, Banff, Canada, 2005.

[24] Karthik Lakshminarayanan, Venkata N. Padmanabhan, and Jitendra padhye. Bandwidth estimation in broadband access networks. In *Proceedings of the Internet Measurement Conference*, Taormina, Italy, 2004.