April 2011

An Analysis on End-To-End Inference Methods based On Packet Probing in Network

J.Benitha Christinal

Karunya University Coimbatore, Tamil Nadu, India, benithachristinal@karunya.edu.in

V. Vivek

Karunya University Coimbatore, Tamil Nadu, India, vivek@karunya.edu.in

Follow this and additional works at: <https://www.interscience.in/ijcns>

Part of the [Computer Engineering Commons](#), and the [Systems and Communications Commons](#)

Recommended Citation

Christinal, J.Benitha and Vivek, V. (2011) "An Analysis on End-To-End Inference Methods based On Packet Probing in Network," *International Journal of Communication Networks and Security*: Vol. 1 : Iss. 1 , Article 9.

DOI: 10.47893/IJCNS.2011.1007

Available at: <https://www.interscience.in/ijcns/vol1/iss1/9>

This Article is brought to you for free and open access by the Interscience Journals at Interscience Research Network. It has been accepted for inclusion in International Journal of Communication Networks and Security by an authorized editor of Interscience Research Network. For more information, please contact sritampatnaik@gmail.com.

An Analysis on End-To-End Inference Methods based On Packet Probing in Network

J.Benitha Christinal, V.Vivek

Department of Computer Science and Engineering, Karunya University

Coimbatore, Tamil Nadu, India

benithachristinal@karunya.edu.in, vivek@karunya.edu.in

Abstract— *Current Internet is a massive, distributed network which continues to grow in size as globalization takes major role in everyone's life like e-commerce, social networking and related activities grow. The heterogeneous and largely unregulated structure of the Internet renders tasks such as optimized service provision, rate limiting certain classes of applications (e.g. peer-to-peer), provide bandwidth guarantee for certain applications, avoiding shared congestion in flows are increasingly challenging tasks. The problem is complicated by the fact that one cannot rely on the cooperation of individual servers and routers to aid in the collection of network traffic measurements vital for these tasks. Hence we go for network monitoring and inference methods based on packet probing in the network. This paper presents an analysis of different inference methods for network characteristics to deal with shared congestion, packet forwarding priority, network tomography and evaluates each methodology based on packet loss rate and delay variance.*

Keywords— *Inference, Packet forwarding priority, packet loss, congestion control, Network tomography*

I.INTRODUCTION

The Internet was designed with no gatekeepers over new content or services. A lightweight but enforceable neutrality rule is needed to ensure that the internet continues to thrive.

—Vint Cerf

Internet is a massive, distributed network which takes major role in our day-to-day life. As the network grows, the internet has evolved very rapidly in a largely unregulated and open environment. The lack of centralized control and the heterogeneous nature of the internet lead to a very important problem: mapping network connectivity, bandwidth, congestion and performance functions. Wide varieties of network characteristics and internet maps have been produced using existing networking tools such as ping and trace route. Information on these tools, along with a collection of interesting internet mapping projects are found in CAIDA [1] and Network tools[2]. The mapping techniques described in the reference

above, usually provide only a partial picture of the internet and network characteristics.

In this paper we present the Analysis on End-To-End Inference for various Methods for shared congestion, packet forwarding priority, network tomography, measuring services based on Packet Probing in Network is presented here. This paper discusses the various inference methods to measure network characteristics. Rest of the paper is organized as follows. Section II contains a generalized summary of various techniques and Section III give brief description of different end to end inference methods for analysing Network Characteristics that have been taken for study. Section III gives a comparative analysis of various inference based on certain parameters. We conclude in Section IV analysing the network characteristics using various inference methods.

II.GENERAL SURVEY

Inference and prediction of network conditions is of fundamental importance to a range of network-aware applications. We classify and survey these research efforts.

One widely adopted strategy is to mine the data collected by network internal resources, such as Border Gateway Protocol routing tables, to generate performance reports [3, 4, 5]. This approach is best applied over long-time scales to produce aggregated analyses such as Internet data sources and analysis reports, but does not lend itself well to providing answers to the fine grained issues we propose here.

Another approach is statistical inference of network internal characteristics based on end to end measurements of point to point traffic [6, 7, 8, 9]. We adopt this general approach because information is gathered at the appropriate granularity. These approaches can be further classified as active approaches [10], which introduce additional probe traffic into the network, and passive approaches, which make inferences only from existing network traffic. The benefit of the former approach is flexibility: one can make measurements at those locations and times which are most valuable. While the benefit of the latter approach is that no additional bandwidth and network resources are consumed solely for the purpose of data collection.

On other dimensions, one can also classify approaches as either receiver-oriented or sender-oriented, depending on where inferences are made and multicast driven or unicast driven, depending on the model used to transmit probe traffic. These are the most common environment under which all inferences made are studied. The general survey gives the idea about various approaches and methods the analysis is made to infer the network characteristics, network conditions of the internet. In next section we have done an literature survey analysis network characteristics like congestion control inferring shared resources [11, 12, 13], network tomography inferring link level performance and topology information [14, 15, 16] and packet forwarding prioritization inferring network QoS and packet scheduling [17, 18]. From the above end to end network inference methods; few are selectively analyzed in detail in this literature

I. LITERATURE SURVEY

A. Robust identification of shared losses using end-to-end unicast probes [19]

Khaled Harf and Azer Bestavros in their paper titled Robust identification of shared losses using end-to-end unicast probes [19], explain method deals with current internet transport protocols make end-to-end measurements and maintain per-connection state to regulate the use of shared network resources. When two or more such connections share a common endpoint, there is an opportunity to correlate the end-to-end measurements made by these protocols to better diagnose and control the use of shared resources. This paper has developed packet-pair probing technique to determine whether a pair of connections experience shared congestion. Packet-Pair Probing is one of the essential techniques in construction of the use of “packet-pair” techniques, to determine bottleneck bandwidth on a network path. In this work, it uses a packet pair probe to a pair of different receivers to introduce loss and delay correlation.

Estimation of Network Parameters Using End-to-End Measurements (Bayesian Approach) this paper proposes an analytical technique for the robust determination of both loss and bottleneck equivalence for pairs of unicast connections emanating from the same server. It is mainly based on end-to-end loss information available at the server as a result of passive monitoring or of active probing. The two connections sharing common common end point is shown in figure 1 at node 2.

Fig. 1 node 1 to node 2 sharing connection [19]

The above scenario Fig. 1 with single server, which has active connections to two distinct clients, both experiencing steady-state packet loss rate $\bullet$. the path from server to client form a tree, which from server's perspective consists of sequence of shared links followed by sequence of disjoint links, in which the shared portion of the sequence may be empty

1. *Loss sharing*: for these two connections, determines if the incidence of packet loss on the shared portion of the tree is at least $\bullet/k$, for a fixed constant $k > 1$
2. *Bottleneck Equivalence*: for these two connections, determines if the incidence of shared loss is greater than the incidence of disjoint loss.

In this paper, a technique for determining whether a pair of connections emanating from the same node experience shared losses for unicast probes has been presented.

A. Detecting shared congestion of flows via end-to-end measurement [20]

Dan Rubenstein and Jim Kurose in their paper titled Detecting shared congestion of flows via end-to-end measurement [20], presents a technique based on loss or delay observations at end-hosts to infer using Poisson probing whether or not two flows experiencing congestion are congested at the same network resources. It validates these techniques via queuing analysis. Current Internet congestion control protocols operate independently on a per-flow basis. A key technical issue underlying both of these scenarios is the ability to detect whether two “flows” whether individual unicast sessions, or different senders within a single multicast session share a common resource bottleneck. In this paper, it addresses the fundamental issue of detecting shared points of congestion among flows. Informally, the point of congestion (POC) for two flows is the same when the same set of resources (e.g., routers) are dropping or excessively delaying packets from both flows due to backup and/or overflowing of queues.

It presents the technique that operates on an end-to-end basis and use only end-system observations to detect whether or not a pair of flows experiences a common POC. The POC for a flow is the set of locations (routers) at which the flow's packets are lost or experience excessive queuing delay. It says it is *testing* two flows when it is trying to identify whether or not they have the same POC. For conciseness, it say that two flows *share congestion* if their POCs are identical, and that flows *do not share congestion* if the intersection of their POCs is empty. The insight is to construct a measure of correlation between flows and a measure of correlation within a flow with the following property: the measure within the flow is greater than the measure within a flow if and only if the flows share the same POC. We call this method of identifying whether or not two flows share a POC a *comparison test*.

The techniques for detecting whether or: not pair of flows share congestion is based on two fundamental observations of internet congestion:

- Losses or delay experienced by two packets passing through the same POC exhibit some degree of correlation. However, in general, the degree of correlation decreases as the time between the packets' transmission is increased.
- The losses or delays experienced by two packets that do not share the same POC will exhibit little or no correlation.

Thus in this paper a technique has been proposed that, via end-to-end measurement, we are able to accurately detect whether or not two flows share the same points of congestion within the network

B. Internet Tomography[21]

Mark Coates and Alfred Hero in their paper titled Internet Tomography [20], deals with the problem of identifying topology and inferring link-level performance parameters such as packet drop rate or delay variance using only end-to-end measurements. This inference is commonly referred to as network tomography. The heterogeneous and largely unregulated structure of the Internet renders tasks such as dynamic routing, optimized service provision, service-level verification, and detection of anomalous/malicious behaviour increasingly challenging tasks. End users who cannot directly access the network links must use end-to-end measurements in order to infer the variables of interest of a given set of links, which requires solving a system of equations that relate these measurement outcomes with these variables. As this system of equations usually does not have a unique solution, current methods use the unrealistic assumption that all links have the same prior probability of being congested.

Increasingly, network operators do not directly operate computers on their network, yet are responsible for assessing network vulnerabilities to ensure compliance with policies about information disclosure, and tracking services that affect provisioning. Thus, with decentralized network management, service discovery becomes an important part of maintaining and protecting computer networks.

It explores two approaches to service discovery: active probing and passive monitoring. Active probing finds all services currently on the network, except services temporarily unavailable or hidden by firewalls; however, it is often too invasive, especially if used across administrative boundaries. Passive monitoring can find transient services, but miss services that are idle. It compares the accuracy of passive and active approaches to service discovery and show that they are complimentary, highlighting the need for multiple active scans coupled with long-duration passive monitoring. It finds passive monitoring is well suited for quickly finding popular services, finding servers responsible for 99% of incoming connections within minutes. Active scanning is better suited to rapidly finding all servers, which is important for vulnerability detection—one scan finds 98% of services in two hours, missing only a handful. External scans are an unexpected ally to passive monitoring, speeding service discovery by the equivalent of 9-15 days of additional observation.

This paper has provided an overview of the large scale inference and tomography in communication networks by using probing schemes and inference methods.

C. Multiple Source, Multiple Destination Network Tomography[22]

Michael Rabbat and Robert Nowak in their paper titled Multiple Source, Multiple Destination Network Tomography[21], presents a study of the multiple source, multiple destination network tomography problem. Using multiple sources in the context of network tomography, It is possible to identify segments within a network shared by the paths connecting multiple sources and destinations. This information may be useful for identifying potential bottlenecks. Sharing statistics between sources may also be useful for optimizing the use of network resources when transferring large amounts of data. Additionally, in some cases it is possible to fuse information gleaned from multiple sources to get a more accurate and refined network characterization. The majority of work in network tomography has revolved on active probing from a single source. Also, it is typical to focus on either identifying the topology, or estimating link-level performance parameters in which case it is assumed that the topology is known. This paper presents a multiple source active measurement procedure and a statistical framework enabling the joint characterization of topology and link-level performance.

Jointly solving for performance parameters and topology leverages on the close coupling between link-level characteristics, routes derived from the network topology, and end-to-end measurements. Inference and characterization of network properties using active end-to-end measurements is a challenging problem. Because the participating hosts are distributed across the network it is not practical to assume that they can be precisely synchronized. Additionally, labels which apply globally cannot be assigned to internal nodes by topology identification techniques employing end-to-end measurements. In general, internal nodes are only inferred relative to the single source from which measurements are made. Thus, the problem of identifying a multiple source topology amounts to more than just matching nodes with the same label.

This paper focuses on the multiple source, multiple destination network tomography problem of characterizing the topology and performance on links connecting a collection of sources and destinations. The contributions are as follows.

- 1) It is shown that the general network tomography problem can be decomposed into a set of smaller components, each involving just two sources and two destinations and easily extend the results to more general multiple source, multiple destination networks.
- 2) It identifies a dichotomy of possible two-source, two-destination topologies based on the model order of their representations.
- 3) A novel multiple-source probing algorithm is presented for determining the model order of an unknown two-source, two-destination topology.
- 4) A flexible decision-theoretic framework is developed enabling the joint characterization of topology and internal performance.
- 5) The efficacy and accuracy of the probing algorithm and statistical framework are evaluated through simulation.

Multiple source topologies can be decomposed into 2-by-2 networks, thus by solving the 2-by-2 problem it have essentially solved the M-by-N problem. The possible 2-by-2 networks can further be broken down into shared and non-shared classes based on their model order (number of links and nodes). There are two main reasons it is interested in this dichotomy. If the topology is shared then measurements can be combined from both sources to achieve reduced variance estimates of linklevel parameters on the downstream links. Additionally, when the topology is shared then we have more information about topology (namely some information about the placement of joining points) than we would have if each source had actively probed without collaborating. Packet arrival order is determined at the first shared queue. This was the basis of the multiple source probing algorithms. Main highlights of the algorithm include the fact that precise synchronization is not required, either multicast or unicast packets can be used, and no more packets are required than would have been used if the sources probed without collaborating even though we know more at the end of the day. Because the algorithm is founded on a

principle directly related to topology, namely that the arrival order of packets is determined at the joining point – the algorithm is robust to cross-traffic and can operate effectively under a variety of conditions.

This paper has provided a probing algorithm for multiple source, multiple destination tomography in networks by using multiple source probing schemes and inference methods.

E. POPI: A User-level Tool for Inferring Router Packet Forwarding Priority[23]

Guohan Lu, Yan Chen and Stefan Birrer in their paper titled A User-level Tool for Inferring Router Packet Forwarding Priority[22], In this paper, it presents an end-to-end approach for packet forwarding priority inference by measuring the loss rate difference of different packet types and its associated tool, POPI. This tool can be used by the enterprises or end-users to discover whether their traffic are treated differently by the ISPs, and whether the ISPs has fulfilled the contracts between them and the users.

Packet forwarding prioritization (PFP) Packet forwarding prioritization (PFP) in routers is one of the mechanisms commonly available to network operators. PFP can have a significant impact on the accuracy of network measurements, the performance of applications and the effectiveness of network troubleshooting procedures. Despite its potential impacts, no information on PFP settings is readily available to end users. In this paper, it presents an end-to-end approach for PFP inference and its associated tool, POPI. POPI enables users to discover such network policies through measurements of packet losses of different packet types. **Inferring Packet-Forwarding Priority**

There may be several candidate metrics to infer packet forwarding priority, such as packet loss, delay or out-of-order events. In this paper, it only uses packet loss as the inference metric because it is the most direct consequence of a priority configuration. It do not rely on packet delay measurements since they may fail to reveal the priorities experienced by packets, as low-priority packets may simply be dropped under congestion without having experienced significant increases in queuing delays. It do not use packet reordering as the metric since certain priority setting mechanisms such as Policing may not generate out-of-order events at all. PFP in routers are set in a per-interface basis. Prioritization of packets does not become evident until the associated link (or a sub link for a traffic class) is saturated, at which point the configured router will begin to drop packets based on its settings. This simple observation defines the basis of the approach used in POPI: *In order to reveal packet-forwarding priorities, one needs to saturate the path available bandwidth for a given class to produce loss rates difference among different classes.* Assuming the existence of a PFP mechanism in routers such an approach will

succeed at uncovering priority settings in routers along a path if the available bandwidth for the controlled class is lower than the bottleneck available bandwidth of the path.

Fig. 2 A burst consists of $nr \times k$ packets [22]

For every burst Fig. 2, loss rate ranks are computed by first sorting packet types in ascending order according to their packet loss rates in that burst and then assigning ranks in order, i.e. the packet type with the largest loss rate has rank 1, the one with the second largest loss rate has rank 2, and etc on. Similar to packet loss rates, due to randomness of packet losses, the ranks of different packet types are like random arrangements over the all bursts when the packet types are treated equally. Every packet burst can be regarded as an observation. Identifying whether there is consistent difference among k ranks over n observations is a well-known statistical problem called *problem of n rankings*. Classic non-parametric solutions such as the Friedman test can find whether there is consistent difference, but they do not make partitions among packet types. Therefore, we proposed to use Average Normalized Ranks (ANR) to group packet types when there is consistent difference. The ANR is the average of the ranks for a packet type over all bursts.

Performance Analysis on Priority Group Partitioning

It first simulate many sets of random rank values (given the number of priority groups and packet types) that satisfy the following two conditions:

- 1) For all packet type i belongs to priority group G_u , and for all packet type j belongs to priority group G_v , the loss rate rank $r_i > r_j$ when G_u has higher priority than G_v .
- 2) For packet types within the same priority group, their ranks are randomly permuted in each burst in order to simulate the effects of random losses.

In this paper, it has demonstrated that POPI, an end-to-end priority inference tool, is able to accurately infer the router's packet forwarding priority using loss statistics.

I. ANALYSIS

Packet-level measurement is now critical to many aspects of broadband networking, for example for guaranteeing service level agreements, facilitating measurement-based admission control algorithms and performing network tomography. Because it is often impossible to measure the entire data passing across a network, the most widely used method of measurement works by injecting probe packets. The probes provide samples

of the packet loss and delay, and from these samples the loss and delay performance of the traffic as a whole can be deduced. However, measuring performance like this is prone to errors. Using packet probing method we have analysed many network characteristics and comparison of inference methods is made.

Parameters used for Comparison

The main parameters we considered for the analysis on End-To-End Inference Methods Based on Packet Probing in Network are Probing methods, Technique to Evaluate, Packet Loss Statistics, Packet Delay statistics, probing rate, Queuing Discipline and Topology.

Packet Probing: Packet probing is an important Internet measurement technique, supporting the investigation of packet delay, path, and loss. Current packet probing techniques use Internet Protocols such as the Internet Control Message Protocol (ICMP), the User Datagram Protocol (UDP), and the Transmission Control Protocol (TCP) to infer network Characteristics.

Technique to Evaluate: It is a procedure used to accomplish a specific activity or task. The network characteristics like shared congestion, Congestion control, Network tomography, Internet tomography, tomography using multiple sources and multiple destinations, packet forwarding prioritization can be evaluated using mentioned techniques.

Packet Loss Statistics: Packet loss occurs when one or more packets of data travelling across a computer network fail to reach their destination. To understand packet loss, it is first necessary to know that information is sent over the Internet in packets. These packets contain all the information needed for the sending computer to communicate the desired information to the destination. In many cases, these packets arrive without any problems. When problems do occur, packet loss can take place. It is one of the most frustrating aspects of digital communications. Here we specify where packet loss exactly occurs in network during congestion.

Packet Delay statistics: In computer networking, packet delay variation is the difference in end-to-end delay between selected packets in a flow with any lost packets being ignored. The effect is sometimes, incorrectly, referred to as jitter. The delay is specified from the start of the packet being transmitted at the source to the end of the packet being received at the destination. Here we specify effect of packet delay in network and how it affects the network.

Probing rate: we use probe packets to measure the packet level performance (e.g. loss, delay); for example whether it is best to probe at a uniform rate, high, or to send probes according to some renewal process, such as a Poisson process. This can be inferred using probe rate.

Queuing Discipline: Queuing Discipline represents the way the queue is organised (rules of inserting and removing customers to/from the queue). Queues are identified by a

handle <major number; minor number>, where the minor number is zero for queues. Handles are used to associate classes to queuing disciplines. Queuing disciplines and classes are tied to one another. The presence of classes and their semantics are fundamental properties of the queuing disciplines. There are many queues like FIFO, CBQ, RED, Drop Tail etc., which are used for Queuing is analysed.

Topology: Network topology is the layout pattern of interconnections of the various elements ([links](#), [nodes](#), etc.) of a [computer network](#). Topology can be considered as a virtual shape or structure of a network. This shape does not correspond to the actual physical design of the devices on the computer network. Any particular network topology is determined only by the graphical mapping of the configuration of physical and/or logical connections between nodes.

These are the parameters which we used for the analysis and comparison of various techniques which we used to infer the network characters. Thus we have compared papers based on shared congestion in unicast environment, shared congestion and congestion control on multicast environment, network tomography to infer topology information and loss statistics, internet tomography, tomography with multiple sources and multiple destinations. The comparisons of the characteristics of all these inference methods are given in TABLE I. We have analysed and studied many papers on End-User level inference to study network characteristics. Then we selected five papers which have similar approach, techniques or network statistics which is used to analysis the network and internal parameters.

TABLE I COMPARISON OF END-USER LEVEL INFERENCE METHODS

Techniques	Shared congestion using unicasting	Shared congestion using multicasting	Internet Tomography	Network Tomography	Packet Forwarding Prioritization
probing	Packet pair probing	Poisson probing	Multicast probing	Semi-Randomized probing	Link probing
Topology	Multicast tree topology	Y topology and inverted Y topology	Cluster topology	Multiple sources, Multiple Destination topology	Dumbbell topology
Packet Loss Statistics	Common source, common destination	Point of congestion (router)	When cross traffic is more	Link level loss	Forcing packets to drop
Packet Delay Statistics	Packet reorder	Packet reorder	Due to cross traffic	Path delay, link delay	No cross traffic so there is no variability in delay
Probing Rate	High	High	Normal	Slow	Faster
Queuing Discipline	Drop Tail	RED	Drop Tail/RED	Drop Tail/FIFO	CBQ/ Priority Queuing

V.CONCLUSION

In this paper, we studied an analysis of different inference methods for network characteristics to deal with shared congestion, packet forwarding priority, network tomography and evaluate each methodology based on packet loss rate and delay variance. We have analysed the strengths and weaknesses of various inference methods and evaluated the techniques based on the packet loss and packet delay statistics. Our evaluation shows the inference methods at End-user level will help the users and network administrators to know network characteristics which are private at router level through various approaches.

REFERENCES

- [1] CAIDA: Cooperative Association for Internet Data Analysis <http://www.caida.org/Tools/>.
- [2] Networking tools and monitoring devices. <http://network-tools.com/>
- [3] Felix: Independent Monitoring for Network Survivability. <http://belore.Com/pub/mwg/felix/index.html>.
- [4] IPMA : Internet Performance Measurement and Analysis. <http://www.Merit.edu/ipma>.
- [5] Mtrace : Tracing multicast path between a source and a receiver. <http://Xerox.com/pub/netsearch/ipmulti>
- [6] R.Caceres, N. G. Duffield, S. B. Moon, and D. Towsley. Inference of Internal Loss Rates in the Mbone. In *IEEE Global Internet (GlobeCom)*, Rio de Janeiro, Brazil, 1999.
- [7] V. Padmanabhan. Optimizing Data Dissemination and Transport in the Internet. Presented at the BU/NSF Workshop on Internet Measurement, Instrumentation and Characterization, September 1999.
- [8] S. Ratnasamy and S.McCanne. Inference of multicast routing trees and Bottleneck bandwidths using end-to-end measurements. In *Proceedings Of IEEE INFOCOM'99*, pages 353-60, March 1999.
- [9] M. Yajnik, S. Moon, J. Kurose, and D. Towsley. Measurement and Modelling of the temporal dependence in packet Loss. In *Proceedings Of IEEE INFOCOM '99*, pages 345-52, March 1999.
- [10] H. Balakrishnan, H. Rahul, and S. Seshan. An Integrated Congestion *SIGCOMM'99*, Cambridge, MA, September 1999.
- [11] V. Padmanabhan. Coordinated Congestion Management and Bandwidth Sharing for Heterogeneous Data Streams. In *Proceedings of NOSSDAV'99*, Basking Ridge, NJ, June 1999.
- [12] L. Gautier, C. Diot, and J. Kurose. End-to-end Transmission Control for Multiparty Interactive Applications in the Internet. In *Proceedings of IEEE INFOCOM'99*, New York, NY, March 1999.
- [13] J. Byers, M. Luby, and M. Mitzenmacher. Accessing Multiple Mirror Sites in Parallel: Using Tornado Codes to Speed Up Downloads. In *Proceedings of IEEE INFOCOM'99*, New York, NY, March 1999.
- [14] A. Adams, T. Bu, R. Caceres, N.G. Duffield, T. Friedman, J. Iroowitz, F. Lo Presti, S.B. Moon, V. Paxson, and D. Towsley. "The Use of End-to-End Multicast Measurements for Characterizing Internal Network Behavior", *IEEE Communications Magazine*, May 2000.
- [15] T. Bu, N.G. Duffield, F. Lo Presti, and D. Towsley. "Network Tomography on General Topologies". *UMass CMPSCI Technique Report*.
- [16] M.J. Coates and R. Nowak. "Network Delay Distribution Inference from End-to-end Unicast Measurement," *Proc. Of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, May 2001.
- [17] M. Aron, P. Druschel, and W. Zwaenepoel. Cluster reserves: A Mechanism for resource management in cluster-based network servers. In *Proceedings of ACM SIGMETRICS 2000*, June 2000.
- [18] J. Bennett and H. Zhang. WFWQ: Worst-case FairWeighted Fair Queueing. In *Proceedings of IEEE INFOCOM '96*, San Francisco, CA, March 1996.
- [19] K. Harfoush, A. Bestavros, and J. Byers, "Robust identification of Shared losses using end-to-end unicast probes," in *Proc. IEEE ICNP*, 2000.
- [20] D. Rubenstein, J. Kurose, and D. Towsley, "Detecting shared Congestion of flows via end-to-end measurement," *IEEE/ACM Transactions on Networking*, vol. 10, no. 3, pp. 381-395, 2002.
- [21] A. Coates, A. Hero III, R. Nowak, and B. Yu, "Internet tomography," *IEEE Signal Processing Magazine*, vol. 19, no. 3, pp. 47-65, 2002.
- [22] M. Rabbat, R. Nowak, and M. Coates, "Multiple source, multiple destination network tomography," in *Proc. IEEE INFOCOM*, 2004.
- [23] G. Lu, Y. Chen, S. Birrer, F. E. Bustamante, C. Y. Cheung, and X. Li, POPI: A User-level Tool for Inferring Router Packet Forwarding Priority, in *Proc. IEEE INFOCOM*, 2010.