Computer Networks 55 (2011) 1052–1068

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

Estimating network link characteristics using packet-pair dispersion: A discrete-time queueing theoretic analysis

Bikash Kumar Dey [a], D. Manjunath [a,*], Supriyo Chakraborty [b,1]

[a] Indian Institute of Technology, Bombay, India
[b] University of California, Los Angeles, CA, United States

ARTICLE INFO

Article history:
Received 10 March 2010
Accepted 12 May 2010
Available online 16 June 2010
Responsible Editor: I.F. Akyildiz

Keywords:
Packet-pair probing
Network tomography
Capacity estimation

ABSTRACT

We present a queueing theoretic analysis of packet-dispersion based probing. The links are modeled as independent discrete-time queues with i.i.d. arrivals. We first derive an algorithm to obtain the (joint) distribution of the separation between the probes at the destination(s) for a given distribution of the spacing at the input. The parameter estimates of the arrival processes are obtained as the minimizer of a cost function between the empirical and calculated distributions. We also carry out extensive simulations and numerical experiments on the model to study the performance of the estimation algorithm for some non stationary arrival process. We find that the estimations work fairly well for two queues in series and for multicast. We also identify issues related to packet-dispersion based parameter estimation when there are multiple queues.

1. Introduction and background

Network tomography is the estimation of detailed statistics of parameters of interest from aggregate or end-to-end measurements of some measurable quantity. The term was first coined in [1]. In this paper, we are interested in the analysis of link bandwidth estimation tomography problems. Traffic matrix estimation [2–6] and link delay estimation [7–9] are other well studied network tomography problems.

There are two commonly used path bandwidth metrics. The *bottleneck bandwidth* is the minimum of the transmission rates on the links in the path. The *available bandwidth* is the portion of the bottleneck bandwidth not used by competing traffic and depends on the traffic load at the inputs to the links. Many tools have been developed to measure the bottleneck and available bandwidths, e.g., `pathchar` [10], `clink` [11], `pathload` [12], and `pathChirp` [13].

To the best of our knowledge, all bandwidth measurement tools are based on either 'packet-pair' or 'packet-dispersion' techniques. In the packet-pair technique [14] two back-to-back packets of equal length are transmitted by the source. The ratio of the probe packet length to the separation between them at the receiver is the service rate for the packet at the bottleneck link on the path where the links are rate-servers, i.e., under a fluid model for the cross traffic. `pathchar` and `clink` are based on this technique and estimate the bottleneck capacity of a network path. `clink` also estimates the transmission rates on each of the links on a path.

The packet-dispersion, or packet-spacing, technique is variation of the packet-pair in which a source transmits a number of probe packets with a predetermined separation. The samples of the separation at the receiver are then used in the bandwidth estimators for the path. `pathload` [12] and `pathChirp` [13] are examples of tools that use the packet-spacing technique to measure the available bandwidth. See [15] for an excellent survey of the different packet-spacing techniques used in the measurement tools and [16] for an experimental comparison of the tools. A non cooperative version that does not

* Corresponding author. Tel.: +91 22 25767427.
E-mail address: dmanju@ee.iitb.ac.in (D. Manjunath).

[1] Work was done while at IIT Bombay.

1389-1286/$ - see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.comnet.2010.05.008

require a measurement-enabled receiver is described in [17].

Packet-pair and packet-dispersion techniques have an important advantage in that they do not require the sender and the receiver to be synchronized in time. This makes them important in practice and motivates the theoretical study of the technique. In this paper, we study bandwidth estimation using packet-pair dispersion in a queueing theoretic setting. Towards this, we define a queueing model for the path of a packet, albeit with simplifications that are amenable to analysis. Many statistical studies of network traffic processes at links show that the packet arrival processes are quite complex and defy simple characterization. A queueing theoretic analysis that accounts for these behaviours is also very messy. Further, as has been pointed out in [18,19] the manner in which the random processes that affect the packet-dispersion interact is also not amenable to easy analysis. Of course, our simplifications do not address many practical situations but provide insight into the performance of packet-dispersion based estimators. We will see that the simplified model helps us obtain several insights into performance of the packet-dispersion based bandwidth estimation tools.

There have been a few theoretical models developed for packet-dispersion based bandwidth probing techniques. We now delineate previous theoretical and theory based literature from the work reported in this paper. In [18,19] it is shown that the output dispersion of probe pair samples the sum of three correlated random processes that are derived from the traffic arrival process. The asymptotics of these and the probing process are also analysed. The measurement methods of some bandwidth estimation tools are then related to the analytical models developed. These results are then extended to multihop paths in [20]. The packet dispersion on a IEEE 802.11 wireless network is analysed in [21]. Fluid model based estimators are developed in [22,23]. Also, see [24] for a discussion of the effectiveness of packet-pair based techniques. In [25], the network is assumed to be a time-invariant, min-plus system that has an unknown service curve. This service curve estimation method is developed for different probing techniques.

Some of the above works provide interesting theoretical insights into the probing process, but our work is more closely related to that in [26–29]. In [26], the bottleneck link is modeled as an M/D/1 queue, and the expected dispersion is calculated using the transient analysis. The method of moments is then used to estimate the arrival rate. A similar method is followed in [27] where the more general M/G/1 queue is considered and Takacs' integro-differential equation is used to obtain the distribution of the dispersion. The method of [26] is extended in [28] to use feedback to obtain a more robust estimator. In [29], a probe train is inserted and the actual delay of each of the probe packets is assumed to be known at the output of the path. The delay sequence is assumed to form a Markov chain whose transition probability matrix is estimated. An inversion is defined to estimate the characteristics of the arrival process from the transition probabilities of the delay Markov chain.

Another class of analysis of packet-dispersion based algorithms is via a statistical analysis. See [30,31] for excellent statistical studies performed on experimental data.

1.1. Problem overview and contributions of the paper

We analyse packet dispersion on unicast paths and multicast trees. The links on a path, or a tree, are modeled as discrete-time queues. The arrival process to each queue are assumed to be independent. This is a reasonable assumption because we can expect that there is sufficient multiplexing at the bottleneck links so that arrival processes, and hence the delays experienced by the packets, on them can also be assumed to be independent. Further, the packet arrivals to each queue in a slot are assumed to be i.i.d. The service time of a packet is equal to the slot length.

The model is not as restrictive as it appears. To account for packet lengths, we can discretise time sufficiently finely, so that the probability that there are more than two packet arrivals in a slot is negligible. Expressing the packet length as an integral multiple of the slot length, the number of packet arrivals in a slot will correspond to the length of the packet arriving in the slot. For multiple queues on a path, a slot on each link is assumed to contain the same number of bits. This implies that the duration of a slot could be different on different links. Thus the packet-pair dispersion at the output of a link needs to be suitably scaled to obtain the dispersion at the input to the next queues. This is possible if the transmission rates on the links are known, or estimated using tools such as `clink` [11]. Although our analysis easily incorporates this change of scale, for simplicity of exposition we assume that all links have the same transmission rate.

In the system that we consider, packet-pair probes with predetermined separation between them are injected at a source. These probe packets traverse discrete-time queues on a path or on a multicast tree. For a single queue and for a given separation between the probes at the input, we first derive the conditional distribution of the separation between the probes at the output of the queue in terms of the distribution of the arrival process. The key result here is the joint distribution of the number of arrivals to the queue and the number of departures from the queue between the slots in which the probes are injected. Using this, the distribution of the output separation is obtained for any given distribution of the input separation. Since the input separation distribution is known, (e.g., we assume it to be fixed) this is applied recursively on the path to obtain the separation distribution at the output. For multicast trees, the joint distribution of the separations at all the outputs is obtained. We first develop our results by assuming that the probe packets are served after all the packets that arrive in their slot, i.e., the probe packets have the least priority among the packets that arrive in their slot. The case of the probe packet having the same priority (or any arbitrary fixed priority) as the other packets that arrive in the slot is also presented.

The application of the preceding analysis for bandwidth estimation is as follows. We assume that the model for the process is known. This reduces the traffic estimation problem to that of parameter estimation. We obtain the samples of the separation between the probe packets at the output nodes of the path or the multicast tree. A possible parameter estimator is the minimizer of a suitable distance