Theory and Practice of Cross-Traffic Estimation

Sridhar Machiraju
UC Berkeley

Darryl Veitch
CUBIN, University of Melbourne

François Baccelli
INRIA-ENS, Ecole Normale Supérieure

Antonio Nucci
Narus

Jean Bolot
Sprint ATL

ABSTRACT

Active probing heuristics are usually based on queuing systems. However, a rigorous probabilistic treatment of probing methods has been lacking. For instance, it is not known even in principle, what can and cannot be measured in general, nor the true limitations of existing methods. We provide a probabilistic treatment for the measurement of cross traffic in the 1-hop case. We derive inversion formulae for the cross traffic process, and explain their fundamental limits, using an intuitive geometric framework.

Categories and Subject Descriptors

C.4 [**Performance of Systems**]: Measurement Techniques, Modeling Techniques, Performance Attributes.

General Terms

Algorithms, Measurement, Performance.

Keywords

Cross-Traffic Estimation, Active Probing.

1. INTRODUCTION

Active probing has become one of the main ways in which the performance of IP networks such as the Internet are measured. In active probing, a stream of *probe packets* are injected from a source at the edge of the network, and collected at cooperating receiver(s). From the known packet sizes, and the measured timestamps of emission and reception, information on both static network parameters and traffic conditions can be inferred.

The literature initially focused on link bandwidth estimation and in particular that of the smallest link, the *bottleneck bandwidth*. More recently, estimation of *available bandwidth* along a path, which is a function not only of the physical network but also of all the *cross traffic* impacting on the probes, has attracted considerable interest. Far less attention has been paid to the idea of using probing to measure queuing processes in routers, or properties of the cross traffic itself (see however [2] and the references therein). In fact, there are deeper underlying issues of system identifiability which require investigation. There is currently no consensus on what can and cannot in fact be measured.

In this paper we make what we believe to be one of the first steps toward the understanding of the in-principle potential and limitations of active probing methods. We do this by looking in some detail into the problem of measuring the distribution of the cross traffic process through the histories probes accumulate in traversing a hop. Since the probes interact with the cross traffic via queuing systems, it can be viewed as an 'inverse queuing problem'. It necessitates considering, in a joint fashion, the statistics of both the queuing process and that of the arriving traffic, and is very different from the traditional questions studied in queuing theory. We address this question in a simple yet well motivated 1-hop setting.

SIGMETRICS'05, June 6–10, 2005, Banff, Alberta, Canada.
ACM 1-59593-022-1/05/0006.

2. DESCRIBING THE SYSTEM

We consider a simplified problem consisting of a single hop, whose FIFO single server queue has a deterministic service rate μ, and an infinite buffer. We take the probes to be of a constant size of p bytes, corresponding to $x = p/\mu > 0$ seconds of workload. Let $\{T_n\}$ and $\{T'_n\}$ be the sequence of arrival and departure times respectively of the probes to the queue. $\{T_n\}$ is deterministic (i.e., periodic) with fixed inter-arrival time t. However, our approach is general enough to be applied with other probing streams such as packet pairs. The raw data of a probing experiment are the departure times, or equivalently, the end-to-end hop *delays* $\{D_n = T'_n - T_n\}$. We describe the input traffic in terms of a *random measure* A, whereby the workload (measured in seconds after dividing by μ) arriving to the queue in a time interval I is denoted by the random variable $A(I)$. In this way we include point or continuous arrivals in a unified and general framework.

Our aim is to recover as much information as possible about the cross traffic described by A using the measured delays. For this to be feasible, we assume that A is stationary and that the sequence of end-to-end probe delays is stationary and ergodic. We also assume that the sequence of stochastic processes $\{A([it, it+v)), 0 \leq v \leq t\}_{i\in\mathbf{N}}$ is i.i.d. Using standard queuing theory for FIFO queues (for example see [1]), the equation describing probe delays can be written as

$$T'_{n+1} = x + \big[\big(T'_n + A([T_n, T_{n+1}))\big) \vee \sup_{v\in[T_n,T_{n+1}]} (v + A([v, T_{n+1}))\big]$$

where $x \vee y$ denotes the maximum of x and y. The left hand argument of $\vee$ dominates iff the probes are in the same busy period.

Subtracting T_{n+1} from both sides of the equation above, a recursive relationship emerges for the delay series $D_n = T'_n - T_n$. We work with residual delays $\{R_n\}$, where $R_n = D_n - x \geq 0$, i.e., the excess delay above the minimum value of x, the probe service time. Considering pairs of consecutive delays R_i and R_{i+1} as samples of random variables R, S we can derive,

$$S = (x + R + C) \vee B \quad \text{where} \tag{1}$$

$$C = A([0, t)) - t, \tag{2}$$

$$B = \sup_{v\in[0,t]}(A([v, t)) - (t - v)). \tag{3}$$

Note that B and C are functionals of the cross traffic over the interval $[0, t)$ *only*. We can interpret C as the net work that arrives in the t time units between consecutive probes. and it takes values in $(-t, \infty)$. Thus, C gives information on the *integral* of the cross traffic over a typical probe inter-arrival, whereas B gives some information on the *peak*. The following important relationships hold by definition:

$$0 \leq B \quad \text{and} \quad C \leq B \leq C + t \tag{4}$$

3. IN-PRINCIPLE INVERSION

To proceed with the inversion, we henceforth assume that all variables, including time, are discrete. This assumption is not essential, as the discretization can be made as fine as we wish. We denote the discrete density and the 2-dimensional cumulative distribution function (CDF) of (B, C) respectively by

$$h(k, l) = P(B = k, C = l), \tag{5}$$

$$H(k, l) = P(B \leq k, C \leq l), \tag{6}$$

in the domain of definition $k \geq 0$, $l \geq -t$. The relationships listed in Equation (4) imply that $h(k, l) = 0$ outside of the diagonally oriented 'feasible strip' defined by (see Figure 1)

$$\text{feasible strip:} \quad (k, l) : k - t \leq l \leq k, \quad k \geq 0. \tag{7}$$

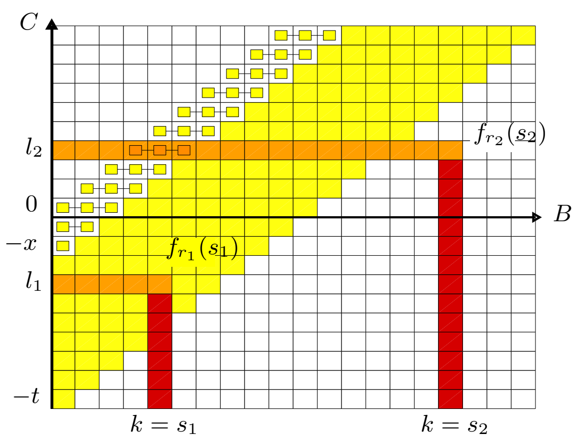

Figure 1: The domain $\{k, l\}$ where the joint density $h(k, l)$ of (B, C) vanishes is shown as white. The support of (B, C) is the strip shown as the light colored band. Two angle sets corresponding to different $f_r(s)$ probabilities are shown. The region where $x + 1 = 3$ masses are connected is the *exclusion zone* where individual h values cannot be resolved.

3.1 Inversion Expressions

Let $f_r(s) = P(S = s|R = r)$ and $F_r(s) = P(S \leq s|R = r)$ be its CDF. We can write down the first conditional probability by accounting the cases when either the left or the right hand arguments in Equation (1) equals s. Due to the independence of R from (B, C), this can be simply written as

$$\begin{aligned} f_r(s) = {} & P(B \leq s, C = s - r - x) \\ & + P(B = s, C \leq s - r - x - 1) \qquad (8) \\ = {} & H(s, s - r - x) - H(s - 1, s - r - x - 1). \qquad (9) \end{aligned}$$

This probability corresponds to a sum of $h(k, l)$ over an 'angle shaped' set, with corner at $(k^*, l^*) = (s, s - r - x)$. Two examples of angle sets are illustrated in Figure 1. A particular observation of $S = s$ given $R = r$ corresponds to a $(B, C) = (k, l)$ value, which, although unobservable, must lie inside the angle set defined by (r, s). We see that the available information concerning $h(k, l)$ comes in the form of the probabilities, given by $f_r(s)$ for different observed (r, s), of falling into different angle sets.

Consider then the possible locations of the 'corners' of these angle sets. For a fixed r, as s is increased the corresponding corner values $(k^*, l^*) = (s, s - r - x)$ move upward, tracing out a line parallel to the main diagonal. As r decreases these diagonals translate upward. But the highest of these, corresponding to $r = 0$, is not the upper boundary of the strip, but lies below it on the line $l = k - x$. Since the corners of angle sets do not lie above the line $l = k - x$, the connected masses shown in Figure 1 cannot be individually resolved using the conditional probabilities. We refer to these h values as belonging to an *exclusion* zone of size $x + 1$ where x is the probe service time. The rest of the h values can be determined, in principle, using Equation (9). We state, without proof, the following result for $k - l - x > 0$:

$$\begin{aligned} h(k, l) = {} & F_{k-l-x}(k) + F_{k-l-x}(k - 1) - \\ & F_{k-l-x+1}(k) - F_{k-l-x-1}(k - 1). \qquad (10) \end{aligned}$$

3.2 Practical Estimation

Define $c(l)$ and $C(l)$ to be the PDF and CDF of C respectively. Figure 1 shows that $c(l)$ can be approximated as:

$$\sum_{i=0}^{\infty} h(i, l) = c(l) \approx c(k, l) = \sum_{i=0}^{k} h(i, l). \tag{11}$$

In the (B, C) plane, $c(k, l)$ is a horizontal bar at level $C = l$ from $B = 0$ till $B = k$. Moreover, for $k \geq l + t$, $c(l)$ is exactly $c(k, l)$. For instance, in Figure 1, $c(l_2) = c(s_2, l_2)$. We also observe that $f_{r_2}(s_2)$ is equal to $c(s_2, l_2)$. This is because the darker vertical bar of $f_{r_2}(s_2)$ is in the white region where h is canonically 0. In general, $f_{r_1}(s_1)$ only approximates $c(s_1, l_1)$.

Hence, $c(l)$ can be approximated by $f_r(l + x)$ and $C(l)$ by $F_r(l + x)$. This approximation is exact for $r \geq t - x$. Estimation accuracy can be improved by using a formula that combines data of multiple r values. For example,

$$c(l) = P(S - R = u|R \geq r) = g_r(u) \tag{12}$$

Equations (10) and (12) can be used to estimate h and c using observed delays. However, the ambiguity due to the exclusion zone prevents us from inverting the PDF of B.

4. CONCLUSIONS AND FUTURE WORK

We use an intuitive geometric framework to provide insight into the in-principle potential of probing for the 1-hop case, under suitable assumptions. We are currently fine-tuning practical estimators based on the foundations set forth in this paper. Our evaluation methodology is to use simulations, traces of cross-traffic from routers and "live" experiments involving the injection of probes on a path that has real cross-traffic on it. We intend to use the insights from these evaluations to address probe stream design too.

5. REFERENCES

[1] F. Baccelli and P. Bremaud. *Elements of Queueing Theory*. Springer Verlag, Applications of Mathematics, second edition, 2003.

[2] A. Pásztor. *Accurate Active Measurement in the Internet and its Applications*. PhD thesis, University of Melbourne, Victoria 3010, Australia, 2003.